

A $ZPP^{NP[1]}$ Lifting Theorem

Thomas Watson*

July 15, 2020

Abstract

The complexity class $ZPP^{NP[1]}$ (corresponding to zero-error randomized algorithms with access to one NP oracle query) is known to have a number of curious properties. We further explore this class in the settings of time complexity, query complexity, and communication complexity.

- For starters, we provide a new characterization: $ZPP^{NP[1]}$ equals the restriction of $BPP^{NP[1]}$ where the algorithm is only allowed to err when it forgoes the opportunity to make an NP oracle query.
- Using the above characterization, we prove a *query-to-communication lifting theorem*, which translates any $ZPP^{NP[1]}$ decision tree lower bound for a function f into a $ZPP^{NP[1]}$ communication lower bound for a two-party version of f .
- As an application, we use the above lifting theorem to prove that the $ZPP^{NP[1]}$ communication lower bound technique introduced by Göös, Pitassi, and Watson (ICALP 2016) is not tight. We also provide a “primal” characterization of this lower bound technique as a complexity class.

1 Introduction

Query-to-communication lifting is a paradigm for proving lower bounds in communication complexity [KN97, Juk12, RY20] using lower bounds in query complexity (a.k.a. decision tree complexity) [Ver99, BdW02, Juk12]. This technique has yielded a wide array of applications, including lower bounds for the Clique vs. Independent Set communication game and the related Alon–Saks–Seymour conjecture in graph theory [Göös15, BDHT17], separations between communication complexity and partition number [GPW18a, AKK16, GJPW18, ABB⁺16, GPW17], lower bounds for monotone circuits, monotone span programs, and proof complexity [RM99, BEGJ00, Joh01, GP18, dNV16, RPRC16, PR17, GGKS18, PR18, GKRS19], new and unified proofs of quantum communication lower bounds [She11] and of separations between randomized and quantum communication complexity [GPW17, AA18, ABK16], lower bounds for LP and SDP relaxations of CSPs [CLRS16, LRS15, KMR17], separations between communication complexity classes [BVd07, Kla11, GLM⁺16, GPW18b, GKPW19, BCH⁺17], lower bounds for finding Nash equilibria [RW16, BR17, GR18], and data structure lower bounds [CKLM18].

The basic format of the technique is a two-step approach in which a relatively simple problem-specific argument is combined with fairly heavy-duty general-purpose machinery for handling communication protocols. More specifically:

*Department of Computer Science, University of Memphis. Supported by NSF grant CCF-1657377.

Class	Reference
P	[RM99, GPW18a]
NP	[GLM ⁺ 16, Göö15]
BPP	[GPW17]
P^{NP}	[GKPW19]
$\text{ZPP}^{\text{NP}[1]}$	Theorem 2
SBP	[GLM ⁺ 16]
AWPP	[She11]
PostBPP	[GLM ⁺ 16]
PP	[She11]

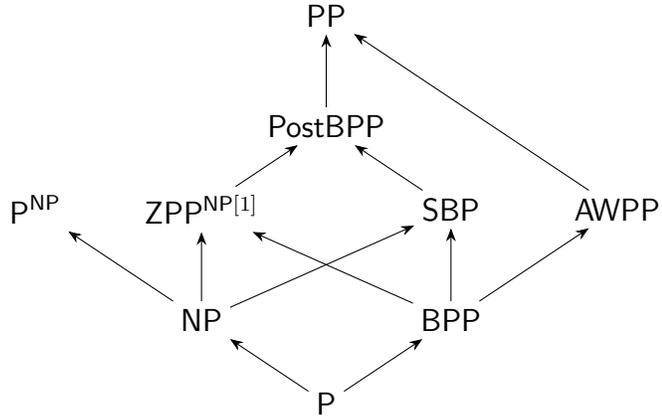


Figure 1: Classes with a known query-to-communication lifting theorem. $\mathcal{C}_1 \rightarrow \mathcal{C}_2$ denotes $\mathcal{C}_1 \subseteq \mathcal{C}_2$.

- (1) Capture the combinatorial core of the desired communication complexity lower bound by proving an analogous query complexity lower bound.
- (2) Apply a *lifting theorem* showing that the query complexity of any boolean function f is essentially the same as the communication complexity of a two-party version of f .

The availability of a lifting theorem greatly eases the burden on the lower bound prover, since query lower bounds are generally much easier to prove than communication lower bounds.

The lifting theorem is with respect to a particular model of computation: deterministic, randomized, nondeterministic, and so on; it is convenient to associate these models with their corresponding classical time-bounded complexity classes: P, BPP, NP, and so on. This idea has led to an ongoing project: prove lifting theorems for the query/communication analogues of various classical complexity classes. Figure 1 shows the main classes for which a lifting theorem is known, along with primary references. Even the less well-known classes sometimes correspond to standard measures in the query/communication settings; e.g., AWPP corresponds to approximate polynomial degree in query complexity and to log of approximate rank in communication complexity. Some notable classes for which a lifting theorem is not known include BQP, UP, and MA. Proving a lifting theorem for AM would be a breakthrough, as it is notoriously open to prove any strong AM-type communication lower bound for an explicit function, but is trivial to do so in the query complexity setting.

Our central contribution is a lifting theorem for the slightly exotic class $\text{ZPP}^{\text{NP}[1]}$, which corresponds to randomized algorithms that can make one call to an NP oracle, output the correct answer with probability $\geq 3/4$, and output \perp with the remaining probability. This model is interesting partly because it has so many curious properties, one of which is that it is robust with respect to the success probability threshold: by [CP08], the success probability can be efficiently amplified as long as it is $> 1/2$ (which is nontrivial since the standard method for amplification would use multiple independent trials, resulting in multiple NP oracle queries). In terms of relations to other classes, $\text{ZPP}^{\text{NP}[1]}$ contains BPP [CC06] and is contained in S_2P [CC06] and in PostBPP (a.k.a. BPP_{path}) [GPW18b]. If we generalized $\text{ZPP}^{\text{NP}[1]}$ to allow success probability slightly $< 1/2$, or to allow two nonadaptive NP oracle calls, either way the class would contain $\text{AM} \cap \text{coAM}$, and hence proving explicit lower bounds for the communication version would yield breakthrough AM communication lower bounds; in this sense, $\text{ZPP}^{\text{NP}[1]}$ is just shy of the communication lower bound

frontier. $\text{ZPP}^{\text{NP}[1]}$ also shows up frequently in the literature on the “two queries problem” [Tri10].

Our starting point is to uncover another curious property of $\text{ZPP}^{\text{NP}[1]}$: we prove it is equivalent (in time, query, and communication complexities) to a new model we dub $\text{CautiousBPP}^{\text{NP}[1]}$, which corresponds to randomized algorithms that can make one call to an NP oracle, output the correct answer with probability $\geq 3/4$, and are only allowed to err when they choose not to call the NP oracle. This equivalence plays a crucial role in our proof of the lifting theorem for $\text{ZPP}^{\text{NP}[1]}$.

Once we have the lifting theorem, the natural application domain is to prove new $\text{ZPP}^{\text{NP}[1]}$ -type communication lower bounds. [GPW18b] developed a technique for proving such lower bounds, and we use our lifting theorem to derive new separations, which imply that the technique from [GPW18b] is not tight. This is analogous to the main application from [GKPW19], in which a P^{NP} lifting theorem was used to show that the P^{NP} -type communication lower bound technique from [IW10, PSS14] is not tight. For context, we note that certain other communication complexity classes have similar lower bound techniques that *are* tight; e.g., the discrepancy bound captures PP communication [Kla07], and the corruption bound captures SBP communication [GW16]. So for what class is the lower bound technique from [GPW18b] tight, if not $\text{ZPP}^{\text{NP}[1]}$? We also answer this question. The class did not have a standard name, but it turns out to have a reasonably simple definition.

1.1 Statement of results

We formally define $\text{ZPP}^{\text{NP}[1]}$ and $\text{CautiousBPP}^{\text{NP}[1]}$ and their query/communication analogues in §2. For any model \mathcal{C} (such as $\text{ZPP}^{\text{NP}[1]}$ or $\text{CautiousBPP}^{\text{NP}[1]}$) we use \mathcal{C} for the polynomial time complexity class, \mathcal{C}^{dt} and \mathcal{C}^{cc} for the polylog query and communication complexity classes, and $\mathcal{C}^{\text{dt}}(f)$ and $\mathcal{C}^{\text{cc}}(F)$ for the corresponding query and communication complexities of a partial function $f: \{0, 1\}^n \rightarrow \{0, 1\}$ and a partial two-party function $F: \{0, 1\}^n \times \{0, 1\}^n \rightarrow \{0, 1\}$ (we also consider F 's where Alice and Bob have unequal but polynomially-related input lengths). We use $\tilde{\Theta}$ to hide polylog(n) factors. We prove the following characterization in §3.

Theorem 1.

- (i) $\text{ZPP}^{\text{NP}[1]} = \text{CautiousBPP}^{\text{NP}[1]}$.
- (ii) $\text{ZPP}^{\text{NP}[1]\text{dt}}(f) = \tilde{\Theta}(\text{CautiousBPP}^{\text{NP}[1]\text{dt}}(f))$ for all f .
- (iii) $\text{ZPP}^{\text{NP}[1]\text{cc}}(F) = \tilde{\Theta}(\text{CautiousBPP}^{\text{NP}[1]\text{cc}}(F))$ for all F .

We now prepare to state the lifting theorem. For $f: \{0, 1\}^n \rightarrow \{0, 1\}$ (called the *outer function*) and $g: \mathcal{X} \times \mathcal{Y} \rightarrow \{0, 1\}$ (called the *gadget*), their composition $f \circ g^n: \mathcal{X}^n \times \mathcal{Y}^n \rightarrow \{0, 1\}$ is the two-party function where Alice gets $x = (x_1, \dots, x_n) \in \mathcal{X}^n$, Bob gets $y = (y_1, \dots, y_n) \in \mathcal{Y}^n$, and the goal is to evaluate $(f \circ g^n)(x, y) := f(g(x_1, y_1), \dots, g(x_n, y_n))$. Note that any deterministic (P-type) decision tree for f can be turned into a deterministic protocol for $f \circ g^n$ where Alice and Bob communicate to evaluate $g(x_i, y_i)$ whenever the decision tree queries the i^{th} input bit of f . A similar thing can be done in other models besides deterministic. The essence of a lifting theorem is to go in the other direction: convert a protocol for $f \circ g^n$ into a comparable-cost decision tree for f . In other words, if g is sufficiently complicated, then it hides the input bits to f so well that a communication protocol cannot do any better than just running a decision tree for f .

We use the *index gadget* $\text{IND}_m: [m] \times \{0, 1\}^m \rightarrow \{0, 1\}$ mapping $(x, y) \mapsto y_x$, where m is a sufficiently large polynomial in n . This gadget has previously been used for the P, BPP, and P^{NP} lifting theorems. (In some cases, lifting theorems with simpler gadgets are known—such as

the “inner product mod 2” gadget on $O(\log n)$ bits [GLM⁺16, WYY17, CKLM19]—but for many applications the index gadget is fine.)

Theorem 2. *Let $m = m(n) := n^C$ for a large enough constant C . For every $f: \{0, 1\}^n \rightarrow \{0, 1\}$,*

- (i) $\text{ZPP}^{\text{NP}[1]\text{cc}}(f \circ \text{IND}_m^n) = \tilde{\Theta}(\text{ZPP}^{\text{NP}[1]\text{dt}}(f))$,
- (ii) $\text{CautiousBPP}^{\text{NP}[1]\text{cc}}(f \circ \text{IND}_m^n) = \Theta(\text{CautiousBPP}^{\text{NP}[1]\text{dt}}(f) \cdot \log n)$.

Note that part (i) of Theorem 2 is a corollary of part (ii), since by Theorem 1,

$$\begin{aligned} \text{ZPP}^{\text{NP}[1]\text{cc}}(f \circ \text{IND}_m^n) &= \tilde{\Theta}(\text{CautiousBPP}^{\text{NP}[1]\text{cc}}(f \circ \text{IND}_m^n)) \\ &= \tilde{\Theta}(\text{CautiousBPP}^{\text{NP}[1]\text{dt}}(f)) = \tilde{\Theta}(\text{ZPP}^{\text{NP}[1]\text{dt}}(f)). \end{aligned}$$

We are not aware of a way to prove part (i) directly, without going through Theorem 1. To prove part (ii) (in §4), we combine tools and techniques from the proofs of lifting theorems for BPP [GPW17], NP [GLM⁺16, Göö15], and P^{NP} [GKPW19], along with some new technical contributions.

Two of the main results in [GPW18b] are $\text{MA}^{\text{cc}} \not\subseteq \text{ZPP}^{\text{NP}[1]\text{cc}}$ and $\text{US}^{\text{cc}} \not\subseteq \text{ZPP}^{\text{NP}[1]\text{cc}}$, where MA and US are the classes associated with “Merlin–Arthur games” and “unique witnesses” respectively (more precise definitions are deferred to §5). The proofs introduced a certain lower bound technique—let us use $\mathcal{B}^{\text{cc}}(F)$ for the largest bound attainable for F using this technique, and \mathcal{B}^{cc} for the class of all F ’s with $\mathcal{B}^{\text{cc}}(F) \leq \text{polylog}(n)$ —and showed that $\text{MA}^{\text{cc}} \not\subseteq \mathcal{B}^{\text{cc}}$, $\text{US}^{\text{cc}} \not\subseteq \mathcal{B}^{\text{cc}}$, and $\text{ZPP}^{\text{NP}[1]\text{cc}} \subseteq \mathcal{B}^{\text{cc}}$. The definition of \mathcal{B}^{cc} is not important for now, but we provide it in §6, where we show that it can be characterized as a more natural complexity class.

Since $\text{ZPP}^{\text{NP}[1]\text{cc}}$ is closed under complement (whereas \mathcal{B}^{cc} is not), we have $\text{ZPP}^{\text{NP}[1]\text{cc}} \subseteq \mathcal{B}^{\text{cc}} \cap \text{co}\mathcal{B}^{\text{cc}}$. A natural question is whether the latter is actually an equality, i.e., whether the lower bound technique of [GPW18b] for $\text{ZPP}^{\text{NP}[1]\text{cc}}$ is tight. Since [GPW18b] observed that $\text{MA}^{\text{cc}}, \text{US}^{\text{cc}} \subseteq \text{co}\mathcal{B}^{\text{cc}}$, we have $\text{MA}^{\text{cc}} \cap \text{coMA}^{\text{cc}}, \text{US}^{\text{cc}} \cap \text{coUS}^{\text{cc}} \subseteq \mathcal{B}^{\text{cc}} \cap \text{co}\mathcal{B}^{\text{cc}}$, and thus the following result (proven in §5 using Theorem 2) answers this question in the negative (in two different ways).

Theorem 3.

- (i) $\text{MA}^{\text{cc}} \cap \text{coMA}^{\text{cc}} \not\subseteq \text{ZPP}^{\text{NP}[1]\text{cc}}$.
- (ii) $\text{US}^{\text{cc}} \cap \text{coUS}^{\text{cc}} \not\subseteq \text{ZPP}^{\text{NP}[1]\text{cc}}$.

We mention that Theorem 3.(ii) answers an open question from [HR90] by implying that $\text{DP}^{\text{cc}} \cap \text{coDP}^{\text{cc}} \not\subseteq \text{P}^{\text{NP}[1]\text{cc}}$ (since $\text{US}^{\text{cc}} \subseteq \text{DP}^{\text{cc}}$; see §6 for the definition of DP^{cc}). The latter separation has subsequently been generalized by [PSW20]. It was also shown in [PSW20] that $\text{DP}^{\text{cc}} \cap \text{coDP}^{\text{cc}} = \text{P}^{\text{NP}[1]\text{cc}}$ when the classes are restricted to contain only total functions, and thus Theorem 3.(ii) is false for total functions. It remains open whether Theorem 3.(i) holds for total functions.

2 Definitions

We set up notation and provide the formal definitions of $\text{ZPP}^{\text{NP}[1]}$ and $\text{CautiousBPP}^{\text{NP}[1]}$. For the query and communication complexity versions, we follow the convention of using the complexity class names as complexity measures. That is, $\mathcal{C}^{\text{dt}}(f)$ denotes the minimum cost of any correct \mathcal{C} -type decision tree for f , and \mathcal{C}^{dt} also denotes the class of families of partial f ’s with $\mathcal{C}^{\text{dt}}(f) \leq \text{polylog}(n)$; similarly, $\mathcal{C}^{\text{cc}}(F)$ denotes the minimum cost of any correct \mathcal{C} -type communication protocol for F ,

and \mathcal{C}^{cc} also denotes the class of families of partial F 's with $\mathcal{C}^{\text{cc}}(F) \leq \text{polylog}(n)$ (assuming Alice and Bob have polynomially-related input lengths).

In the query complexity setting, “query” actually has two meanings for us: a decision tree makes queries to individual input bits, then it forms an NP-type (DNF) oracle query.

We think of a randomized algorithm M as taking a uniformly random string $s \in \{0, 1\}^r$ (for some number of coins r that depends on the input length); we let $M_s(x)$ denote M running on input x with outcome s . Similarly, we think of a randomized (in our case, ZPP^{NP[1]}-type or CautiousBPP^{NP[1]}-type) decision tree T or communication protocol Π as the uniform distribution over a multiset of corresponding non-randomized T_s 's or Π_s 's indexed by $s \in \{0, 1\}^r$; we denote this as $T \sim \{T_s : s \in \{0, 1\}^r\}$ or $\Pi \sim \{\Pi_s : s \in \{0, 1\}^r\}$.

In general, the single NP oracle query could happen in the middle of a computation, but WLOG we assume it happens at the very end. This is because instead of making the oracle query, an algorithm can remember the query, proceed with the rest of the computation under both possible scenarios (the oracle returns 0 or 1), and finally make the query knowing what the output will be in either scenario—which is captured by a little function we call *out*.

We use L to denote both a language and its characteristic function.

2.1 ZPP^{NP[1]}

ZPP^{NP[1]} consists of all languages L for which there is a polynomial-time randomized algorithm M (taking input x and coin tosses $s \in \{0, 1\}^r$) and a language $L' \in \text{NP}$ such that the following hold.

Syntax: The computation of $M_s(x)$ produces an oracle query q and a function *out*: $\{0, 1\} \rightarrow \{0, 1, \perp\}$; the output is then *out*($L'(q)$).

Correctness: The output is always $L(x)$ or \perp , and is $L(x)$ with probability $\geq 3/4$.

We define a ZPP^{NP[1]}-type decision tree T for f on input x as follows.

Syntax: $T \sim \{T_s : s \in \{0, 1\}^r\}$ where each T_s makes queries to the bits of x until it reaches a leaf, which is labeled with a DNF D and a function *out*: $\{0, 1\} \rightarrow \{0, 1, \perp\}$; the output is then *out*($D(x)$).

Correctness: The output is always $f(x)$ or \perp , and is $f(x)$ with probability $\geq 3/4$.

Cost: The maximum height of any T_s , plus the maximum width of any DNF appearing at a leaf.

We define a ZPP^{NP[1]}-type communication protocol Π for F on input (x, y) as follows.

Syntax: $\Pi \sim \{\Pi_s : s \in \{0, 1\}^r\}$ where each Π_s communicates until it reaches a leaf, which is labeled with a multiset of rectangles $\{R^w : w \in \{0, 1\}^k\}$ (for some k) and a function *out*: $\{0, 1\} \rightarrow \{0, 1, \perp\}$; the output is then *out* applied to the indicator of whether $(x, y) \in \bigcup_w R^w$.

Correctness: The output is always $F(x, y)$ or \perp , and is $F(x, y)$ with probability $\geq 3/4$.

Cost: The maximum communication cost of any Π_s , plus the maximum k at any leaf.

A priori, the value $3/4$ seems arbitrary since it is not clear whether ZPP^{NP[1]} is amenable to amplification of the success probability (naively doing repeated trials would increase the number of NP queries). However, [CP08] showed that amplification is actually possible, so we may use any constant $> 1/2$ for the success probability in the definition of ZPP^{NP[1]} (while affecting the measures ZPP^{NP[1]dt}(f) and ZPP^{NP[1]cc}(F) by only constant factors).

2.2 CautiousBPP^{NP[1]}

CautiousBPP^{NP[1]} consists of all languages L for which there is a polynomial-time randomized algorithm M (taking input x and coin tosses $s \in \{0, 1\}^r$) and a language $L' \in \text{NP}$ such that the following hold.

Syntax: The computation of $M_s(x)$ either directly outputs a bit (without invoking the oracle) or produces an oracle query q and a nonconstant function $out: \{0, 1\} \rightarrow \{0, 1\}$; in the latter case the output is then $out(L'(q))$.

Correctness: The output is $L(x)$ with probability $\geq 3/4$, and is $L(x)$ for all s such that $M_s(x)$ makes an oracle query.

We define a CautiousBPP^{NP[1]}-type decision tree T for f on input x as follows.

Syntax: $T \sim \{T_s : s \in \{0, 1\}^r\}$ where each T_s makes queries to the bits of x until it reaches a leaf, which is labeled with either an output bit, or a DNF D and a nonconstant function $out: \{0, 1\} \rightarrow \{0, 1\}$; in the latter case the output is then $out(D(x))$.

Correctness: The output is $f(x)$ with probability $\geq 3/4$, and is $f(x)$ for all s such that $T_s(x)$ makes a DNF query.

Cost: The maximum height of any T_s , plus the maximum width of any DNF appearing at a leaf.

We define a CautiousBPP^{NP[1]}-type communication protocol Π for F on input (x, y) as follows.

Syntax: $\Pi \sim \{\Pi_s : s \in \{0, 1\}^r\}$ where each Π_s communicates until it reaches a leaf, which is labeled with either an output bit, or a multiset of rectangles $\{R^w : w \in \{0, 1\}^k\}$ (for some k) and a nonconstant function $out: \{0, 1\} \rightarrow \{0, 1\}$; in the latter case the output is then out applied to the indicator of whether $(x, y) \in \bigcup_w R^w$.

Correctness: The output is $F(x, y)$ with probability $\geq 3/4$, and is $F(x, y)$ for all s such that $\Pi_s(x, y)$ makes a “union of rectangles” query.

Cost: The maximum communication cost of any Π_s , plus the maximum k at any leaf.

The success probability of any CautiousBPP^{NP[1]}-type computation can be amplified by taking the majority vote of multiple independent trials—except if at least one of those trials results in an NP-type oracle query then (to avoid making multiple oracle queries) we just use the output of one such trial since we know it will be correct. Thus just like for BPP-type computations, success probability $1/2 + \varepsilon$ can be amplified to $1 - \delta$ with a $O(\frac{1}{\varepsilon^2} \log \frac{1}{\delta})$ factor overhead in cost.

3 ZPP^{NP[1]} = CautiousBPP^{NP[1]}

We now prove [Theorem 1](#), starting with part (i). First assume $L \in \text{ZPP}^{\text{NP}[1]}$, witnessed by a randomized algorithm M (taking input x and coin tosses $s \in \{0, 1\}^r$) and $L' \in \text{NP}$. To see that $L \in \text{CautiousBPP}^{\text{NP}[1]}$, consider the following randomized algorithm with oracle access to L' :

1. Sample $s \in \{0, 1\}^r$ and run $M_s(x)$ until it produces q and out .
2. If $out(0) = out(1)$ then output this common bit, or an arbitrary bit if $out(0) = out(1) = \perp$.
3. Else if one of $out(0), out(1)$ is \perp then output whichever is not \perp .
4. Else invoke the oracle on q and output $out(L'(q))$.

Consider any s for which this algorithm outputs the wrong bit: then it did not make an oracle query (since M never outputs the wrong bit), and $M_s(x)$ would have output \perp (because of either line 2, or line 3 with $out(L'(q)) = \perp$ and $out(1 - L'(q)) \neq L(x)$). Hence this algorithm correctly solves L , with error probability at most that of M .

For the converse direction, we generalize the argument from [CC06] that $BPP \subseteq ZPP^{NP[1]}$. Assume $L \in \text{CautiousBPP}^{NP[1]}$, witnessed by a randomized algorithm M (taking input x and coin tosses $s \in \{0, 1\}^r$) and $L' \in \text{NP}$. Assume that this has already been amplified so the error probability is $< 1/4r$ (by the remark at the end of §2.2). For a fixed input x and $b \in \{0, 1\}$, let

$$S_b := \{s \in \{0, 1\}^r : M_s(x) \text{ outputs } b \text{ without invoking the oracle}\}.$$

To see that $L \in ZPP^{NP[1]}$, consider the following randomized algorithm:

1. Sample $s \in \{0, 1\}^r$ and run $M_s(x)$ until it produces either an output b (so $s \in S_b$) or q and out .
2. If it produced q and out then ask the NP oracle for the value of $L'(q)$ and output $out(L'(q))$.
3. Else sample independent strings $s^1, \dots, s^{4r} \in \{0, 1\}^r$ and ask the NP oracle whether $\bigcup_i (S_b \oplus s^i) \neq \{0, 1\}^r$ (i.e., whether there exists an s' such that for every i , $s' \oplus s^i \notin S_b$); output \perp if so and b if not.

Note that this algorithm never outputs the wrong bit: if $s \in S_b$ for $b = 1 - L(x)$, then $|S_b| < 2^r/4r$ so by a union bound, $|\bigcup_i (S_b \oplus s^i)| < 4r \cdot (2^r/4r) = 2^r$ and hence the NP oracle returns 1 on line 3 and the algorithm outputs \perp . For the success probability, consider two cases. If $|S_0 \cup S_1| \leq 2^r/4$, then line 2 executes (guaranteeing correct output) with probability $\geq 3/4$. Otherwise, since $|S_{1-L(x)}| < 2^r/4r$, we must have $|S_{L(x)}| > 2^r/4 - 2^r/4r > 2^r/5$ (we may assume r is at least a large enough constant), so by a union bound over all $s' \in \{0, 1\}^r$, the probability over s^1, \dots, s^{4r} that $\bigcup_i (S_{L(x)} \oplus s^i) \neq \{0, 1\}^r$ is $< 2^r \cdot (4/5)^{4r} \leq (5/6)^r \leq 1/5$. In this latter case, the probability of outputting \perp is

$$\mathbb{P}[b = 1 - L(x)] + \mathbb{P}[\bigcup_i (S_b \oplus s^i) \neq \{0, 1\}^r \mid b = L(x)] \cdot \mathbb{P}[b = L(x)] \leq 1/4r + (1/5) \cdot |S_{L(x)}|/2^r \leq 1/4.$$

In both cases the success probability is $\geq 3/4$.

Parts (ii) and (iii) are proved in the same way as part (i), but we must carefully analyze the cost. Let us summarize the differences. For the $ZPP^{NP[1]} \subseteq \text{CautiousBPP}^{NP[1]}$ direction, exactly the same argument works but using T_s or Π_s in place of M_s , and making the same DNF query or “union of rectangles” query rather than the same NP oracle query on line 4. This shows that $\text{CautiousBPP}^{NP[1]dt}(f) \leq ZPP^{NP[1]dt}(f)$ and $\text{CautiousBPP}^{NP[1]cc}(F) \leq ZPP^{NP[1]cc}(F)$.

Now consider the $\text{CautiousBPP}^{NP[1]} \subseteq ZPP^{NP[1]}$ direction for parts (ii) and (iii). By standard sparsification of the randomness [New91], we may assume T or Π uses only $O(\log n)$ coin tosses (while affecting the success probability by only $\pm o(1)$). Then as noted at the end of §2.2, we may amplify with $O(\log \log n)$ repetitions so r becomes $O(\log n \cdot \log \log n)$ and the error probability becomes $\leq 1/\log^2 n < 1/4r$. As above, we use T_s or Π_s in place of M_s , and make the same DNF query or “union of rectangles” query rather than the same NP oracle query on line 2. For line 3, we note that the predicate $\bigcup_i (S_b \oplus s^i) \neq \{0, 1\}^r$, as a function of the input x or (x, y) , can be computed by nondeterministically guessing s' and running $T_{s' \oplus s^i}(x)$ or $\Pi_{s' \oplus s^i}(x, y)$ for each $i \in [4r]$; this can be expressed as a DNF of width $4r \cdot (\text{cost of amplified } T)$, or as a union of 2^k rectangles

with $k = r + 4r \cdot (\text{cost of amplified } \Pi)$. Thus, the overall cost is $O(r \cdot (\text{cost of amplified } T \text{ or } \Pi)) \leq O((\text{cost of original } T \text{ or } \Pi) \cdot \log n \cdot \log^2 \log n)$. This shows that

$$\begin{aligned} \text{ZPP}^{\text{NP}[1]\text{dt}}(f) &\leq O(\text{CautiousBPP}^{\text{NP}[1]\text{dt}}(f) \cdot \log n \cdot \log^2 \log n), \\ \text{ZPP}^{\text{NP}[1]\text{cc}}(F) &\leq O(\text{CautiousBPP}^{\text{NP}[1]\text{cc}}(F) \cdot \log n \cdot \log^2 \log n). \end{aligned}$$

(A factor of $\log \log n$ can be shaved off here, by amplifying the original algorithm using a random walk of length $O(\log \log n)$ on an expander graph rather than doing the trials fully independently, so r becomes $O(\log n + \log \log n) = O(\log n)$.) This finishes the proof of [Theorem 1](#).

By the way, the above argument implies that the success probability of $\text{ZPP}^{\text{NP}[1]}$ can be amplified from $1/2 + 1/\text{poly}(n)$ to $1 - 1/\text{poly}(n)$ in polynomial time: a $\text{ZPP}^{\text{NP}[1]}$ algorithm can be converted to a $\text{CautiousBPP}^{\text{NP}[1]}$ one with no loss in efficiency or success probability, then the latter can be amplified as remarked at the end of [§ 2.2](#), then this can be converted back to a $\text{ZPP}^{\text{NP}[1]}$ algorithm with $1 - 1/\text{poly}(n)$ success probability by tweaking the parameters in the proof of [Theorem 1](#). [\[CP08\]](#) showed a stronger result—the success probability of $\text{ZPP}^{\text{NP}[1]}$ can actually be amplified to $1 - 1/\exp(n)$ —but our proof is perhaps simpler.

4 Proof of the Lifting Theorem

We now prove [Theorem 2](#). As noted in [§ 1.1](#), we just need to show part *(ii)*. It is straightforward to see that for all f ,

$$\text{CautiousBPP}^{\text{NP}[1]\text{cc}}(f \circ \text{IND}_m^n) \leq O(\text{CautiousBPP}^{\text{NP}[1]\text{dt}}(f) \cdot \log n)$$

since we can have the communication protocol run the optimal decision tree for f , communicating $O(\log n)$ bits to evaluate $\text{IND}_m(x_i, y_i)$ whenever this bit is queried, and if a width- w DNF oracle query is formed then we can convert each of its $\leq n^w$ conjunctions into $\leq m^w$ rectangles, resulting in a “union of rectangles” oracle query that contributes $k = O(w \log n)$ to the cost. Thus, the bulk of the proof is to show that for all f ,

$$\text{CautiousBPP}^{\text{NP}[1]\text{dt}}(f) \leq O(\text{CautiousBPP}^{\text{NP}[1]\text{cc}}(f \circ \text{IND}_m^n) / \log n). \quad (1)$$

In [§ 4.1](#) we provide relevant technical background from the proofs of earlier lifting theorems (mainly the one for BPP [\[GPW17\]](#)). Then in [§ 4.2](#) we describe how to simulate the communication protocol with a decision tree, and in [§ 4.3](#) we prove a key technical lemma.

4.1 Background

Abbreviate $G := \text{IND}_m^n$. We consider deterministic communication protocols on G 's input domain $[m]^n \times (\{0, 1\}^m)^n$, which we view as partitioned into *slices* $G^{-1}(z) = \{(x, y) : G(x, y) = z\}$, one for each $z \in \{0, 1\}^n$. We let $|\Pi|$ denote the worst-case number of bits communicated by a deterministic protocol Π . We use boldface letters for random variables.

Let $\mathbb{H}_\infty(\mathbf{x}) := \min_x \log(1/\mathbb{P}[\mathbf{x} = x])$ denote the usual min-entropy of a random variable \mathbf{x} . Supposing \mathbf{x} is distributed over a set X , we define the *deficiency* of \mathbf{x} as the nonnegative quantity $\mathbb{D}_\infty(\mathbf{x}) := \log |X| - \mathbb{H}_\infty(\mathbf{x})$. A basic property is that if X is a Cartesian product then marginalizing

\mathbf{x} to some coordinates cannot increase the deficiency (as observed in [GPW17]). For a set X we let \mathbf{X} denote a random variable uniformly distributed on X .

The following definition and claim originate in the proof of the lifting theorems for NP, SBP, and PostBPP [GLM⁺16, Göö15]. They describe an invariant that Alice maintains throughout the simulation, and how to restore it (by fixing some coordinates, which will correspond to querying those input bits of f) when it gets violated. A proof of this specific version of the claim can be found in [GKPW19].

Definition 1. A random variable $\mathbf{x} \in [m]^J$ (where $J \subseteq [n]$ is some index set) is called δ -dense if for every nonempty $I \subseteq J$, the coordinates \mathbf{x}_I (marginally distributed over $[m]^I$) have min-entropy rate at least δ , i.e., $\mathbb{H}_\infty(\mathbf{x}_I) \geq \delta \cdot |I| \log m$.

Claim 1 ([GKPW19]). If $A \subseteq [m]^J$ then there exist an $I \subseteq J$ of size $|I| \leq O(\mathbb{D}_\infty(\mathbf{A})/\log n)$ and a nonempty $A' \subseteq A$ such that A' is fixed on I and 0.9-dense on $J \setminus I$.

It is simple to check that all 2^n slices of G 's input domain have the same size, and that the uniform distribution over any slice is marginally nearly-uniform ($o(1)$ -close in statistical distance) on both Alice's input and Bob's input. The following lemma from [GPW17] provides a sufficient condition for similar properties to hold even after we have queried some of the input bits of f .

Definition 2. For a partial assignment $\rho \in \{0, 1, *\}^n$, define its free positions as $\text{free } \rho := \rho^{-1}(*) \subseteq [n]$, and its fixed positions as $\text{fix } \rho := [n] \setminus \text{free } \rho$. A rectangle $X \times Y$ is called ρ -structured if $\mathbf{X}_{\text{free } \rho}$ is 0.9-dense, $\mathbf{X}_{\text{fix } \rho}$ is fixed, and each element of $G(X \times Y) \subseteq \{0, 1\}^n$ is consistent with ρ .

Definition 3. A distribution \mathcal{D}_1 is said to be ε -pointwise-close to a distribution \mathcal{D}_2 if for every outcome, the probability under \mathcal{D}_1 is within a factor $1 \pm \varepsilon$ of the probability under \mathcal{D}_2 . The distributions are said to be ε -close if the statistical (total variation) distance is $\leq \varepsilon$.

Lemma 1 ([GPW17]). Suppose $X \times Y$ is ρ -structured and $\mathbb{D}_\infty(\mathbf{Y}) \leq n^3$. Then:

- (i) For any $z \in \{0, 1\}^n$ consistent with ρ , the uniform distribution on $G^{-1}(z) \cap X \times Y$ (which is nonempty) has both of its marginal distributions $o(1)$ -close to \mathbf{X} and \mathbf{Y} , respectively.
- (ii) $G(\mathbf{X}, \mathbf{Y})$ is $o(1)$ -pointwise-close to the uniform distribution over the set of all z consistent with ρ .

Now we come to the main part of the proof of the BPP lifting theorem from [GPW17]. It shows that, given query access to z , we can approximately sample the transcript that would be generated by a communication protocol on a random input from z 's slice. In fact, this simulation maintains some invariants, which we need to expose (in the ‘‘furthermore’’ part of the lemma) for use in the subsequent ‘‘NP oracle query’’ phase of our simulation.

Definition 4. A deterministic protocol $\bar{\Pi}$ is said to be a refinement of a deterministic protocol Π if they have the same input domain and for every transcript rectangle $X \times Y$ of $\bar{\Pi}$, there exists a transcript rectangle of Π that contains $X \times Y$.

Lemma 2 ([GPW17]). For every deterministic protocol Π on G 's input domain with $|\Pi| \leq n \log m$, there exist a refinement $\bar{\Pi}$ and a randomized decision tree T of cost $O(|\bar{\Pi}|/\log n)$ that

on input $z \in \{0, 1\}^n$ outputs a transcript of $\overline{\Pi}$ or \perp , such that the following two distributions are $o(1)$ -close:

- $\mathbf{t} :=$ output distribution of T on input z ,
- $\mathbf{t}' :=$ transcript generated by $\overline{\Pi}$ when run on a random input $(\mathbf{x}, \mathbf{y}) \sim G^{-1}(z)$.

Furthermore, for every (non- \perp) transcript output by T on input z with positive probability, the associated rectangle $X \times Y$ satisfies:

- (i) $X \times Y$ is ρ -structured, where ρ corresponds to the results of the queries made by T (and is hence consistent with z),
- (ii) $\mathbb{D}_\infty(\mathbf{Y}) \leq n^{2.5}$,
- (iii) $\mathbb{D}_\infty(\mathbf{X}_{\text{free } \rho}) \leq O(|\Pi|)$.

4.2 Simulation

Lemma 3. *Let $X \times Y$ be a ρ -structured rectangle in G 's input domain such that $\mathbb{D}_\infty(\mathbf{Y}) \leq n^{2.5}$. Suppose $\{R^w \subseteq X \times Y : w \in \{0, 1\}^k\}$ is a collection of rectangles whose union covers exactly $G^{-1}(f^{-1}(1)) \cap X \times Y$. Then f can be computed by a width- $O((\mathbb{D}_\infty(\mathbf{X}_{\text{free } \rho}) + k)/\log n)$ DNF on the domain of inputs consistent with ρ .*

Lemma 3 is our key tool for converting the NP^{cc} oracle query to an NP^{dt} oracle query. The proof, which we give in §4.3, combines insights from the lifting theorem proofs for NP [GLM⁺16, G 15] and P^{NP} [GKPW19] with new calculations. For now we use Lemma 3 to argue (1), thus finishing the proof of Theorem 2.

Let $\Pi \sim \{\Pi_s : s \in \{0, 1\}^r\}$ be a $\text{CautiousBPP}^{\text{NP}[1]}$ -type communication protocol for $f \circ G$ (and note WLOG the cost is $\leq n \log m$). Here is a $\text{CautiousBPP}^{\text{NP}[1]}$ -type decision tree for f on input z :

1. Sample $s \in \{0, 1\}^r$ and (eliding the dependence on s) let $\overline{\Pi}$ and T be the refinement and randomized decision tree from Lemma 2 applied to Π_s .
2. Sample T 's coin tosses s' and run $T_{s'}$ on input z until it either outputs \perp (in which case we halt and output an arbitrary bit) or produces a transcript t of $\overline{\Pi}$.
3. Let $X \times Y$ be the rectangle associated with t , and let t^* be the transcript of Π_s whose rectangle contains $X \times Y$.
4. If t^* outputs a bit, then we halt and output the same bit; otherwise let $\{R^w : w \in \{0, 1\}^k\}$ and $\text{out}: \{0, 1\} \rightarrow \{0, 1\}$ be the rectangles and nonconstant function associated with t^* .
5. Since $X \times Y$ satisfies properties (i), (ii), (iii) from Lemma 2, we may apply Lemma 3 to the collection $\{R^w \cap X \times Y : w \in \{0, 1\}^k\}$ (whose union covers exactly $G^{-1}(f^{-1}(\text{out}(1))) \cap X \times Y$ by the correctness of Π), using f if $\text{out}(1) = 1$ or $\neg f$ if $\text{out}(1) = 0$, to obtain a width- $O((|\Pi_s| + k)/\log n)$ DNF D that computes f or $\neg f$ (respectively) on all inputs consistent with ρ .
6. Output $\text{out}(D(z))$.

Since T makes $O(|\Pi_s|/\log n)$ queries and the DNF on line 5 has width $O((|\Pi_s| + k)/\log n)$, the above decision tree indeed has cost $O((\text{cost of } \Pi)/\log n)$. If it reaches line 5 and makes a DNF query, then the output is correct since z is consistent with ρ and hence $\text{out}(D(z)) = f(z)$. For the success probability, call t good if the corresponding t^* either outputs $f(z)$ directly or makes a

“union of rectangles” query, and note that if the above decision tree generates a good t then the output is correct (by the previous sentence). Hence, letting \mathbf{t}, \mathbf{t}' be the $o(1)$ -close random variables from Lemma 2 applied to Π_s (with $(\mathbf{x}, \mathbf{y}) \sim G^{-1}(z)$), we have

$$\begin{aligned}
\mathbb{P}[\text{output is correct}] &\geq \mathbb{E}_s[\mathbb{P}_{s'}[\mathbf{t} \text{ is a good transcript}]] \\
&\geq \mathbb{E}_s[\mathbb{P}_{\mathbf{x}, \mathbf{y}}[\mathbf{t}' \text{ is good}] - o(1)] \\
&= \mathbb{E}_{\mathbf{x}, \mathbf{y}}[\mathbb{P}_s[\mathbf{t}' \text{ is good}]] - o(1) \\
&= \mathbb{E}_{\mathbf{x}, \mathbf{y}}[\mathbb{P}_s[\Pi_s(\mathbf{x}, \mathbf{y}) \text{ outputs } f(z)]] - o(1) \\
&\geq \mathbb{E}_{\mathbf{x}, \mathbf{y}}[3/4] - o(1) \\
&= 3/4 - o(1).
\end{aligned}$$

We conclude that $\text{CautiousBPP}^{\text{NP}[1]\text{dt}}(f) \leq O(\text{CautiousBPP}^{\text{NP}[1]\text{cc}}(f \circ G) / \log n)$.¹

4.3 Forming a DNF

We now prove Lemma 3. Fix any $z \in f^{-1}(1)$ consistent with ρ , and define $J := \text{free } \rho$. We need to show that there exists a width- $O((\mathbb{D}_\infty(\mathbf{X}_J) + k) / \log n)$ conjunction that accepts z but does not accept any input in $f^{-1}(0)$ consistent with ρ . The final DNF is then the disjunction of these conjunctions over all $z \in f^{-1}(1)$ consistent with ρ .

For each rectangle $R^w = X^w \times Y^w$ define the set of *weighty rows* as

$$A^w := \{x \in X^w : |Y_x^w| \geq 2^{nm-n^3}\} \quad \text{where} \quad Y_x^w := \{y \in Y^w : G(x, y) = z\}.$$

Claim 2. *There exists a $w \in \{0, 1\}^k$ such that $|A^w| \geq |X|/2^{k+1}$.*

Proof. Suppose for contradiction this is not the case. First we show that the weighty rows of all the rectangles do not cover too much of the slice $G^{-1}(z) \cap X \times Y$: By Lemma 1.(i) we have

$$\frac{|G^{-1}(z) \cap (\bigcup_w A^w) \times Y|}{|G^{-1}(z) \cap X \times Y|} \leq \frac{|\bigcup_w A^w|}{|X|} + o(1) \leq \frac{2^k \cdot |X|/2^{k+1}}{|X|} + o(1) < 3/4. \quad (2)$$

On the other hand, we show that the non-weighty rows also do not cover too much of the slice (this part does not rely on our contradiction hypothesis): Since the R^w 's cover $G^{-1}(z) \cap X \times Y$ and since $k \leq n \log m$ WLOG, we have

$$|G^{-1}(z) \cap (X \setminus \bigcup_w A^w) \times Y| \leq |\bigcup_{w, x \notin A^w} Y_x^w| \leq 2^k \cdot |X| \cdot 2^{nm-n^3} \leq |X| \cdot 2^{nm-n^{2.9}},$$

and by Lemma 1.(ii) and $\mathbb{D}_\infty(\mathbf{Y}) \leq n^{2.5} \leq n^3$ we have

$$|G^{-1}(z) \cap X \times Y| \geq |X| \cdot |Y| \cdot (1 - o(1))/2^{|J|} \geq |X| \cdot 2^{nm-n^{2.5}} \cdot (1 - o(1))/2^n \geq |X| \cdot 2^{nm-n^{2.6}},$$

¹Let us summarize the fundamental reason we are unable to make this proof work directly for $\text{ZPP}^{\text{NP}[1]}$ (instead of $\text{CautiousBPP}^{\text{NP}[1]}$) without going through Theorem 1. Suppose we reach line 5 with $\text{out}(1) = \perp$ and $\text{out}(0) \neq \perp$. We would like to form a DNF that accepts those z 's consistent with ρ where $G^{-1}(z) \cap X \times Y$ is covered by the union of $\{R^w \cap X \times Y : w \in \{0, 1\}^k\}$ —and then output \perp if the DNF accepts and output $\text{out}(0)$ if it rejects. The issue is that there may be some z 's consistent with ρ such that $f(z) = \text{out}(0)$ but $G^{-1}(z) \cap X \times Y$ is partially covered by the union—even a fairly small coverage might result in the DNF accepting z . This could cause the overall probability of outputting \perp on z to be much higher in the decision tree than in the communication protocol.

and thus

$$\frac{|G^{-1}(z) \cap (X \setminus \bigcup_w A^w) \times Y|}{|G^{-1}(z) \cap X \times Y|} \leq \frac{|X| \cdot 2^{nm-n^{2.9}}}{|X| \cdot 2^{nm-n^{2.6}}} = 2^{n^{2.6}-n^{2.9}} < 1/4. \quad (3)$$

Now (2) and (3) form a contradiction. This proves the claim. \square

Now fix a $w \in \{0, 1\}^k$ such that $|A^w| \geq |X|/2^{k+1} = m^{|J|}/2^{\mathbb{D}_\infty(\mathbf{X}_J)+k+1}$ and hence $\mathbb{D}_\infty(\mathbf{A}_J^w) \leq \mathbb{D}_\infty(\mathbf{X}_J) + k + 1$. Applying Claim 1 to A_J^w , we can obtain an $I \subseteq J$ of size $|I| \leq O((\mathbb{D}_\infty(\mathbf{X}_J) + k)/\log n)$ and a nonempty $A' \subseteq A^w$ such that A' is fixed on $I \cup \text{fix } \rho$ and 0.9-dense on $J \setminus I$. Consider the conjunction that accepts iff the I coordinates of the input equal z_I ; we now argue that this conjunction satisfies the desired properties. It certainly has the right width and accepts z .

Define $\sigma \in \{0, 1, *\}^n$ as the partial assignment that extends ρ by fixing the coordinates in I to z_I . As in [GKPW19], pick any $x' \in A'$ and let $B := Y_{x'}^w$. Then $A' \times B$ is σ -structured (note that for all $(x, y) \in A' \times B$, $G(x, y)_I = G(x', y)_I = z_I$ since $x_I = x'_I$) and $\mathbb{D}_\infty(\mathbf{B}) \leq n^3$, and thus by Lemma 1.(ii), $G(\mathbf{A}', \mathbf{B})$ is $o(1)$ -pointwise-close to the uniform distribution over all strings consistent with σ . In particular, for every z' consistent with σ (i.e., for every z' consistent with ρ that is accepted by the conjunction) there exists an $(x, y) \in A' \times B$ such that $G(x, y) = z'$; since $A' \times B \subseteq R^w \subseteq G^{-1}(f^{-1}(1))$, this implies that $f(z') = 1$. In summary, the conjunction does not accept any input in $f^{-1}(0)$ consistent with ρ . This finishes the proof of Lemma 3.

5 Applications

We prove Theorem 3 in this section. Since $\text{MA}^{\text{cc}} \cap \text{coMA}^{\text{cc}}$, $\text{US}^{\text{cc}} \cap \text{coUS}^{\text{cc}} \subseteq \mathcal{B}^{\text{cc}} \cap \text{co}\mathcal{B}^{\text{cc}}$, Theorem 3 cannot be shown using the lower bound technique from [GPW18b], so we instead prove the analogous separations in query complexity and apply our lifting theorem. We start by defining the query/communication versions of MA and US.

Merlin–Arthur games (MA) are the model where Merlin nondeterministically sends a message to Arthur (comprised of Alice and Bob in the communication setting), who is randomized and decides whether to accept. On a 1-input, there should exist a witness Merlin can send that makes Arthur accept with probability 1, and on a 0-input, Arthur should reject with probability $\geq 1/2$ no matter what Merlin sends. In the query/communication settings, the cost is Merlin’s message length plus Arthur’s query/communication cost.

The US model is like ordinary nondeterminism, except that an input is accepted iff there is *exactly one* witness that leads to acceptance (so, rejection means there are either 0 or ≥ 2 accepted witnesses). In query complexity, the cost is the maximum width of any of the witness conjunctions. In communication complexity, the cost is the log of the number of witness rectangles.

5.1 MA \cap coMA

We now prove Theorem 3.(i). We start with a general technique for proving $\text{CautiousBPP}^{\text{NP}[1]\text{dt}}$ lower bounds. For a bit b , we say a conjunction is b -monochromatic for a partial function f if it rejects all $(1 - b)$ -inputs.

Lemma 4. *Suppose f has no monochromatic conjunction of width $< k$. Then*

$$\text{CautiousBPP}^{\text{NP}[1]\text{dt}}(f) \geq \min(k, \text{BPP}^{\text{dt}}(f)).$$

Proof. If f has a $\text{CautiousBPP}^{\text{NP}[1]}$ -type decision tree of cost $< k$, then this decision tree must never make a DNF query (in which case it is just a BPP-type decision tree, showing that $\text{BPP}^{\text{dt}}(f) \leq \text{CautiousBPP}^{\text{NP}[1]\text{dt}}(f)$). To see this, suppose for contradiction some leaf is labeled with a DNF query D and a function out , and consider the conjunction that accepts the inputs that lead to that leaf and are accepted by an arbitrarily chosen term of D (which WLOG is consistent with the partial assignment leading to the leaf). Then this conjunction has width $< k$ and is $out(1)$ -monochromatic (as any input accepted by it would make the $\text{CautiousBPP}^{\text{NP}[1]}$ -type decision tree output $out(1)$ after making a DNF query, for some outcome of the coin tosses, and hence could not be an $out(0)$ -input). \square

Let $n = 2\ell^2$, and define the partial function $f: \{0, 1\}^n \rightarrow \{0, 1\}$ that interprets its input as a pair of $\ell \times \ell$ boolean matrices (A, B) , such that $f(A, B) = 1$ iff A has an all-1 row and every row of B is at most half 1's, and $f(A, B) = 0$ iff B has an all-1 row and every row of A is at most half 1's. Note that $f \in \text{MA}^{\text{dt}} \cap \text{coMA}^{\text{dt}}$ since an MA-type decision tree can guess a row in A and check that a random bit from that row is 1, and a coMA-type decision tree can guess a row in B and check that a random bit from that row is 1. This upper bound lifts to $f \circ \text{IND}_m^n \in \text{MA}^{\text{cc}} \cap \text{coMA}^{\text{cc}}$. We now show that $f \notin \text{CautiousBPP}^{\text{NP}[1]\text{dt}}$ which, by [Theorems 1 and 2](#), implies that $f \circ \text{IND}_m^n \notin \text{ZPP}^{\text{NP}[1]\text{cc}}$. This will yield [Theorem 3.\(i\)](#).

By [Lemma 4](#), it suffices to show that

- (1) f has no monochromatic conjunction of width $\leq \ell/2$, and
- (2) $\text{BPP}^{\text{dt}}(f) \geq \Omega(\ell)$.

To see (1), consider any conjunction C of width $\leq \ell/2$: Since it does not touch every row of A , and it touches at most half the bits in each row of B , we can construct a 1-input accepted by C by putting all 1's in an untouched row of A , and filling the rest of the matrix entries with 0's (except those whose value is determined by C accepting). Similarly, there must exist a 0-input accepted by C . Thus C is not monochromatic.

For (2), by the minimax principle it suffices to exhibit an input distribution such that every cost- $o(\ell)$ deterministic decision tree succeeds with probability $< 3/4$ over a random input. We define the input distribution by filling a uniformly random one of the 2ℓ rows with 1's, and letting all other entries of (A, B) be 0's. If the decision tree accepts after seeing only 0's, then conditioned on a random 0-input it continues to accept (and hence err) with probability $\geq 1 - o(1)$ (since the all-0's path of the decision tree only touches a $o(1)$ fraction of the rows). Similarly, if it rejects after seeing only 0's, then conditioned on a random 1-input it continues to reject (and hence err) with probability $\geq 1 - o(1)$. In either case, it errs with probability $\geq 1/2 - o(1)$ over an unconditioned random input.

5.2 $\text{US} \cap \text{coUS}$

We now prove [Theorem 3.\(ii\)](#). Let $\text{weight}(\cdot)$ refer to Hamming weight. For even n , define the partial function $f: \{0, 1\}^n \rightarrow \{0, 1\}$ that interprets its input as $(a, b) \in \{0, 1\}^{n/2} \times \{0, 1\}^{n/2}$, such that $f(a, b) = 1$ iff $\text{weight}(a) = 1$ and $\text{weight}(b) \in \{0, 2\}$, and $f(a, b) = 0$ iff $\text{weight}(b) = 1$ and $\text{weight}(a) \in \{0, 2\}$. Note that $f \in \text{US}^{\text{dt}} \cap \text{coUS}^{\text{dt}}$ since a US-type decision tree can guess the location of a 1 in a , and a coUS-type decision tree can guess the location of a 1 in b . This upper bound lifts to $f \circ \text{IND}_m^n \in \text{US}^{\text{cc}} \cap \text{coUS}^{\text{cc}}$. We now show that $f \notin \text{CautiousBPP}^{\text{NP}[1]\text{dt}}$ which, by [Theorems 1 and 2](#), implies that $f \circ \text{IND}_m^n \notin \text{ZPP}^{\text{NP}[1]\text{cc}}$. This will yield [Theorem 3.\(ii\)](#).

Note that [Lemma 4](#) cannot help us here, since this f does not have small monochromatic conjunctions (e.g., a conjunction with two positive literals from a is 0-monochromatic), so we devise a different technique. In fact, we show something stronger than $f \notin \text{CautiousBPP}^{\text{NP}[1]\text{dt}}$. Define $\text{BPP}^{\text{NP}[1]}$ in the natural way (two-sided error, and allowed to err after an NP oracle query is made), and notice that the class may depend on the exact choice of success probability (since the standard method of amplification involves multiple independent trials, which would increase the number of NP oracle queries). Let us use $\text{BPP}_p^{\text{NP}[1]}$ to indicate that the success probability must be $\geq p$ on each input. As $\text{CautiousBPP}^{\text{NP}[1]}$ can be efficiently amplified (see the end of [§2.2](#)), the following lemma implies that $f \notin \text{CautiousBPP}^{\text{NP}[1]\text{dt}}$.

Lemma 5. *For every constant $\varepsilon > 0$, $\text{BPP}_{3/4+\varepsilon}^{\text{NP}[1]\text{dt}}(f) \geq \Omega(n)$.*

Proof. It suffices to show that every cost- $o(n)$ $\text{P}^{\text{NP}[1]}$ -type decision tree succeeds with probability $\leq 3/4 + o(1)$ over the uniform distribution on valid inputs to f . Let v be the leaf reached after seeing only 0's, and say v is labeled with DNF D and function $out: \{0, 1\} \rightarrow \{0, 1\}$. Assume $out(1) = 1$ (the case $out(1) = 0$ is argued similarly). Consider the joint random variables $\mathbf{a}, \mathbf{b}, \mathbf{a}', \mathbf{b}'$ where \mathbf{a} has a unique 1 in a random position, \mathbf{b} is all 0's, \mathbf{a}' is obtained from \mathbf{a} by flipping a random 0 to 1, and \mathbf{b}' is obtained from \mathbf{b} by flipping a random 0 to 1. Note that (\mathbf{a}, \mathbf{b}) is the input distribution conditioned on $\text{weight}(\mathbf{a}) = 1$ and $\text{weight}(\mathbf{b}) = 0$, and $(\mathbf{a}', \mathbf{b}')$ is the input distribution conditioned on $\text{weight}(\mathbf{a}) = 2$ and $\text{weight}(\mathbf{b}) = 1$. We have $\mathbb{P}[(\mathbf{a}, \mathbf{b}) \text{ reaches } v] \geq 1 - o(1)$ and thus $\mathbb{P}[(\mathbf{a}, \mathbf{b}) \text{ reaches } v \text{ and is accepted}] \geq \mathbb{P}[(\mathbf{a}, \mathbf{b}) \text{ is accepted}] - o(1)$. Also, conditioned on any outcome of (\mathbf{a}, \mathbf{b}) that reaches v and is accepted, with probability $\geq 1 - o(1)$ the two flipped bits are not among those read along the path to v and not among those read by an arbitrarily chosen term of D that accepts (\mathbf{a}, \mathbf{b}) , in which case $(\mathbf{a}', \mathbf{b}')$ also reaches v and is accepted. Thus, $\mathbb{P}[(\mathbf{a}', \mathbf{b}') \text{ reaches } v \text{ and is accepted} \mid (\mathbf{a}, \mathbf{b}) \text{ reaches } v \text{ and is accepted}] \geq 1 - o(1)$. Combining these, we get

$$\begin{aligned} \mathbb{P}[(\mathbf{a}', \mathbf{b}') \text{ is accepted}] &\geq \mathbb{P}[(\mathbf{a}', \mathbf{b}') \text{ and } (\mathbf{a}, \mathbf{b}) \text{ both reach } v \text{ and are accepted}] \\ &= \mathbb{P}[(\mathbf{a}', \mathbf{b}') \text{ reaches } v \text{ and is accepted} \mid (\mathbf{a}, \mathbf{b}) \text{ reaches } v \text{ and is accepted}] \\ &\quad \cdot \mathbb{P}[(\mathbf{a}, \mathbf{b}) \text{ reaches } v \text{ and is accepted}] \\ &\geq (1 - o(1)) \cdot (\mathbb{P}[(\mathbf{a}, \mathbf{b}) \text{ is accepted}] - o(1)) \\ &= \mathbb{P}[(\mathbf{a}, \mathbf{b}) \text{ is accepted}] - o(1). \end{aligned}$$

Thus, under the uniform distribution on valid inputs to f ,

$$\begin{aligned} \mathbb{P}[\text{err}] &\geq \mathbb{P}[\text{err} \mid \text{weight}(\mathbf{a}) = 1 \text{ and } \text{weight}(\mathbf{b}) = 0]/4 + \mathbb{P}[\text{err} \mid \text{weight}(\mathbf{a}) = 2 \text{ and } \text{weight}(\mathbf{b}) = 1]/4 \\ &= (\mathbb{P}[(\mathbf{a}, \mathbf{b}) \text{ is rejected}] + \mathbb{P}[(\mathbf{a}', \mathbf{b}') \text{ is accepted}])/4 \\ &= (1 - (\mathbb{P}[(\mathbf{a}, \mathbf{b}) \text{ is accepted}] - \mathbb{P}[(\mathbf{a}', \mathbf{b}') \text{ is accepted}]))/4 \\ &\geq (1 - o(1))/4. \quad \square \end{aligned}$$

We can complement [Lemma 5](#) by noting that $\text{BPP}_{3/4}^{\text{NP}[1]\text{dt}}(f) \leq 2$: With probability $1/4$ each:

- accept iff $\text{weight}(\mathbf{a}) \leq 1$,
- accept iff $\text{weight}(\mathbf{a}) \geq 1$,
- reject iff $\text{weight}(\mathbf{b}) \leq 1$,

- reject iff $\text{weight}(b) \geq 1$.

Hence $\text{BPP}_{3/4}^{\text{NP}[1]\text{dt}} \not\subseteq \text{BPP}_{3/4+\varepsilon}^{\text{NP}[1]\text{dt}}$, which implies that $\text{BPP}_{3/4}^{\text{NP}[1]} \not\subseteq \text{BPP}_{3/4+\varepsilon}^{\text{NP}[1]}$ in a relativized world. Thus, unlike $\text{ZPP}^{\text{NP}[1]}$, $\text{BPP}^{\text{NP}[1]}$ is not generally amenable to efficient amplification; this phenomenon has subsequently been fully explored in [Wat19].

6 Primal Characterization of the Lower Bound Technique

The following is the lower bound technique introduced in [GPW18b].

Lemma 6 ([GPW18b]). *Suppose μ_0 is a distribution over $F^{-1}(0)$, μ_1 is a distribution over $F^{-1}(1)$, and C is a constant such that for every rectangle R ,*

- (i) $\mu_1(R) \leq \delta$ if R is 1-monochromatic (i.e., contains no 0-inputs), and
- (ii) $\mu_0(R) \leq C \cdot \mu_1(R) + \delta$.

Then $\text{ZPP}^{\text{NP}[1]\text{cc}}(F) \geq \Omega(\log(1/\delta))$.

By an argument of [Kla03], the exact value of the constant C does not matter (as long as it is sufficiently large), only affecting the final bound by constant factors. Thus, there exists a constant C such that $\mathcal{B}^{\text{cc}}(F)$ is defined as the maximum of $\lceil \log(1/\delta) \rceil$ over all μ_0, μ_1, δ satisfying the properties of Lemma 6.

Using only assumption (i) of Lemma 6 would give rise to the so-called 1-monochromatic rectangle size bound, which is well-known to be a tight “dual” characterization of $\text{NP}^{\text{cc}}(F)$ [KN97, §2.4]. Using only assumption (ii) would give rise to the corruption bound, which is known to be a tight “dual” characterization of so-called $\text{coSBP}^{\text{cc}}(F)$ [GW16]. Since the latter model is less widely known, we define it here: An SBP algorithm is randomized and must accept 1-inputs with probability $\geq \alpha$, and accept 0-inputs with probability $\leq \alpha/2$, for some $\alpha > 0$ which is an arbitrarily small function of the input size. In the communication complexity setting, this is equivalent to having a multiset of 2^k rectangles (the cost is k) such that each 1-input (0-input) is contained in $\geq \alpha$ ($\leq \alpha/2$) fraction of them. The query complexity setting is similar but with conjunctions instead of rectangles (the cost is the maximum width).

Since \mathcal{B}^{cc} is defined as a certain dual combination of NP^{cc} and coSBP^{cc} (requiring the bounds to hold simultaneously, under the same input distribution), it is natural to wonder whether it is also a “primal” combination of NP^{cc} and coSBP^{cc} . Specifically, an educated guess is that \mathcal{B}^{cc} functions are exactly those that can be expressed as an “or” of an NP^{cc} function and a coSBP^{cc} function. We now confirm this, and then reprove some lemmas from [GPW18b] through the primal lens. We first set up the relevant notation.

For two partial boolean functions F, H , we let $F \cup H$ denote the partial function such that $(F \cup H)^{-1}(1) = F^{-1}(1) \cup H^{-1}(1)$ and $(F \cup H)^{-1}(0) = F^{-1}(0) \cap H^{-1}(0)$, and we let $F \cap H$ denote the partial function such that $(F \cap H)^{-1}(1) = F^{-1}(1) \cap H^{-1}(1)$ and $(F \cap H)^{-1}(0) = F^{-1}(0) \cup H^{-1}(0)$. Note that $F, H, F \cup H, F \cap H$ may all have different domains. For classes \mathcal{C}, \mathcal{D} , we let $\mathcal{C} \cup \mathcal{D}$ denote the class of all partial functions of the form $F \cup H$ for some $F \in \mathcal{C}$ and $H \in \mathcal{D}$, and we let $\mathcal{C} \cap \mathcal{D}$ denote the class of all partial functions of the form $F \cap H$ for some $F \in \mathcal{C}$ and $H \in \mathcal{D}$.

The class DP may be defined as $\text{NP} \cup \text{coNP}$, and we may assume WLOG in this definition that $F, H, F \cap H$ all have the same domain. The latter assumption cannot always be made for combinations of other classes. For example, consider $\text{NP} \cup \text{coSBP}$: on each 0-input, there is no NP

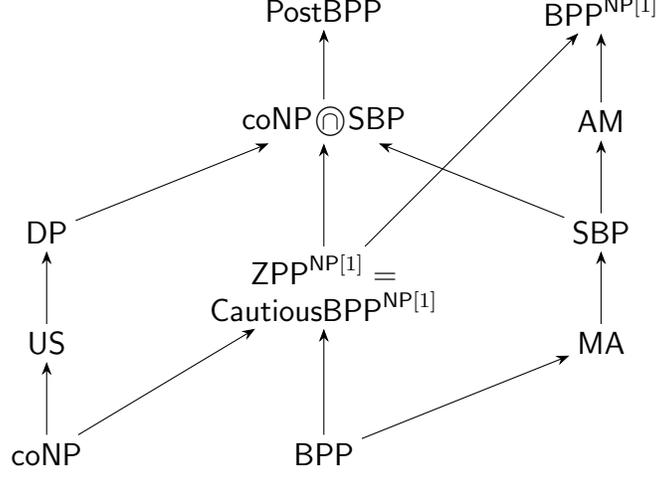


Figure 2: Classes relevant to our applications. $\mathcal{C}_1 \rightarrow \mathcal{C}_2$ denotes $\mathcal{C}_1 \subseteq \mathcal{C}_2$.

witness and the coSBP algorithm accepts with probability $\geq \alpha$, and on each 1-input, either there exists an NP witness or the coSBP algorithm accepts with probability $\leq \alpha/2$. Note that if there exists an NP witness then the coSBP algorithm may accept with illegal probability in the range $(\alpha/2, \alpha)$, in which case the input must not be in the domain of the coSBP partial function. This shows that the NP partial function and the coSBP partial function might not have the same domain as each other.

Theorem 4. $\mathcal{B}^{\text{cc}} = \text{NP}^{\text{cc}} \oplus \text{coSBP}^{\text{cc}}$ and $\text{co}\mathcal{B}^{\text{cc}} = \text{coNP}^{\text{cc}} \oplus \text{SBP}^{\text{cc}}$.

Proof. This is inspired by the argument from [GW16] that coSBP^{cc} captures the corruption bound. A cost- k ($\text{NP} \oplus \text{coSBP}$)-type protocol for F consists of a pair of multisets of rectangles $\{R^w : w \in \{0, 1\}^k\}, \{Q^s : s \in \{0, 1\}^k\}$ and an $\alpha > 0$ such that if $F(x, y) = 0$ then $(x, y) \notin \bigcup_w R^w$ and (x, y) is in $\geq \alpha$ fraction of Q^s 's, and if $F(x, y) = 1$ then either $(x, y) \in \bigcup_w R^w$ or (x, y) is in $\leq \alpha/2$ fraction of Q^s 's. We may assume the latter $\alpha/2$ is actually $\alpha/2^c$ (for any integer $c > 0$), by using the multiset of intersections of all c -tuples of Q^s 's, at the cost of increasing k by a factor of c (and raising α to the power c) [GW16].

First we show that $\mathcal{B}^{\text{cc}}(F) \leq O((\text{NP}^{\text{cc}} \oplus \text{coSBP}^{\text{cc}})(F))$ for all F . Consider an ($\text{NP} \oplus \text{coSBP}$)-type protocol, and assume the coSBP part has already been amplified to have $\alpha/4C$ instead of $\alpha/2$; let k be the new cost, and note that $\alpha \geq 2^{-k}$ WLOG. Consider any μ_0, μ_1 .

Case (i): $\mu_1(\bigcup_w R^w) \geq 2^{-2k}$. Then there exists an R^w (which is 1-monochromatic) such that $\mu_1(R^w) \geq 2^{-3k}$.

Case (ii): $\mu_1(\bigcup_w R^w) < 2^{-2k}$. Let μ_1^* denote μ_1 conditioned on $F^{-1}(1) \setminus \bigcup_w R^w$. For each $s \in \{0, 1\}^k$, let $q_s := \mu_0(Q^s)$ and $r_s := \mu_1^*(Q^s)$. We claim there exists an s such that $q_s \geq \alpha/2$ and $r_s \leq q_s/2C$. Suppose for contradiction the claim is false, and let $S \subseteq \{0, 1\}^k$ be such that for all $s \in S$, $q_s < \alpha/2$, and for all $s \in \bar{S}$, $r_s > q_s/2C$. Then

$$\begin{aligned} \alpha/4C &\geq \mathbb{E}_{(x,y) \sim \mu_1^*} \mathbb{P}_{s \in \{0,1\}^k} [(x,y) \in Q^s] \\ &= \mathbb{E}_{s \in \{0,1\}^k} \mathbb{P}_{(x,y) \sim \mu_1^*} [(x,y) \in Q^s] \end{aligned}$$

$$\begin{aligned}
&= 2^{-k} \sum_{s \in \{0,1\}^k} r_s \\
&\geq 2^{-k} \sum_{s \in \overline{S}} r_s \\
&\geq 2^{-k} \sum_{s \in \overline{S}} q_s / 2C \\
&= 2^{-k} (\sum_{s \in \{0,1\}^k} q_s - \sum_{s \in S} q_s) / 2C \\
&\geq (\mathbb{E}_{s \in \{0,1\}^k} \mathbb{P}_{(x,y) \sim \mu_0} [(x,y) \in Q^s] - 2^{-k} |S| \alpha / 2) / 2C \\
&\geq (\mathbb{E}_{(x,y) \sim \mu_0} \mathbb{P}_{s \in \{0,1\}^k} [(x,y) \in Q^s] - \alpha / 2) / 2C \\
&\geq \alpha / 4C.
\end{aligned}$$

Furthermore, at least one of the inequalities must be strict, which is a contradiction. Now for a fixed such s , we have

$$\begin{aligned}
\mu_0(Q^s) &\geq \mu_0(Q^s) / 2 + \alpha / 4 \geq C \cdot \mu_1^*(Q^s) + \alpha / 4 \geq C \cdot (\mu_1(Q^s) - \mu_1(\bigcup_w R^w)) + \alpha / 4 \\
&\geq C \cdot \mu_1(Q^s) - C \cdot 2^{-2k} + 2^{-k-2} \geq C \cdot \mu_1(Q^s) + 2^{-3k}.
\end{aligned}$$

Combining the two cases shows that if μ_0, μ_1, δ satisfy the properties of [Lemma 6](#), then $\delta \geq 2^{-3k}$ and hence $\log(1/\delta) \leq 3k$.

Now we show that $(\text{NP}^{\text{cc}} \circledast \text{coSBP}^{\text{cc}})(F) \leq O(\mathcal{B}^{\text{cc}}(F) + \log n)$ for all F . Consider any δ such that for all μ_0, μ_1 there exists a rectangle R such that either

- (i) $\mu_1(R) > \delta$ and R is 1-monochromatic, or
- (ii) $\mu_0(R) > C \cdot \mu_1(R) + \delta$.

Our goal is to show that $(\text{NP}^{\text{cc}} \circledast \text{coSBP}^{\text{cc}})(F) \leq O(\log(1/\delta) + \log n)$.

First we record that if (ii) holds above, then there exists a subrectangle $R' \subseteq R$ such that $\mu_0(R') \geq \delta$ and $\mu_1(R') \leq \delta/4$, assuming $C \geq 8$: Let R_1, R_2, \dots be the rows of R sorted in increasing order of $\mu_1(R_i)/\mu_0(R_i)$ and consider the least i such that $\mu_0(R_{\leq i}) \geq \delta$. If $\mu_0(R_{\leq i}) \leq 2\delta$ then $\mu_1(R_{\leq i})/\mu_0(R_{\leq i}) \leq \mu_1(R)/\mu_0(R) < 1/C$ and thus we may take $R' = R_{\leq i}$ since $\mu_1(R_{\leq i}) \leq \mu_0(R_{\leq i})/8 \leq \delta/4$. Otherwise, we have $\mu_0(R_i) > \delta$ and thus we may take $R' = R_i \cap F^{-1}(0)$.

Now let $\mathcal{R}_{(i)}$ be the set of all 1-monochromatic rectangles, and $\mathcal{R}_{(ii)}$ be the set of all other rectangles. Let M be the matrix with rows indexed by inputs (x, y) in the domain of F and columns indexed by rectangles R , such that

$$M_{(x,y),R} := \begin{cases} \delta & \text{if } F(x,y) = 0 \text{ and } R \in \mathcal{R}_{(i)} \\ 1 & \text{if } F(x,y) = 1 \text{ and } R \in \mathcal{R}_{(i)} \text{ and } (x,y) \in R \\ 1 & \text{if } F(x,y) = 0 \text{ and } R \in \mathcal{R}_{(ii)} \text{ and } (x,y) \in R \\ \frac{\delta}{1-\delta/4} & \text{if } F(x,y) = 1 \text{ and } R \in \mathcal{R}_{(ii)} \text{ and } (x,y) \notin R \\ 0 & \text{otherwise} \end{cases}$$

We claim that for every distribution μ over the domain of F , there exists a rectangle R such that $\mathbb{E}[M_{\mu,R}] \geq \delta$. For $b \in \{0,1\}$ let $\pi_b := \mu(F^{-1}(b))$ and $\mu_b := (\mu | F^{-1}(b))$ (or let μ_b be an arbitrary distribution over $F^{-1}(b)$ if $\pi_b = 0$), so $\mu = \pi_0 \mu_0 + \pi_1 \mu_1$. Now by the above, there exists a rectangle R such that either $\mu_1(R) > \delta$ and $R \in \mathcal{R}_{(i)}$, in which case $\mathbb{E}[M_{\mu,R}] = \pi_0 \delta + \pi_1 \mu_1(R) \geq \delta$, or $\mu_0(R) \geq \delta$ (so $R \in \mathcal{R}_{(ii)}$) and $\mu_1(R) \leq \delta/4$, in which case $\mathbb{E}[M_{\mu,R}] = \pi_0 \mu_0(R) + \pi_1 (1 - \mu_1(R)) \frac{\delta}{1-\delta/4} \geq \delta$.

By the minimax theorem, there exists a distribution ν over rectangles such that for every (x, y) in F 's domain, $\mathbb{E}[M_{(x,y),\nu}] \geq \delta$. For $b \in \{(i), (ii)\}$ let $\pi_b := \nu(\mathcal{R}_b)$ and $\nu_b := (\nu|_{\mathcal{R}_b})$, so $\nu = \pi_{(i)}\nu_{(i)} + \pi_{(ii)}\nu_{(ii)}$. Assume for now that $\pi_{(i)}, \pi_{(ii)}$ are both nonzero; we explain later how to handle if one of them is zero. Also, let $\nu_b(x, y) := \mathbb{P}_{R \sim \nu_b}[(x, y) \in R]$.

If $F(x, y) = 0$, we have $\delta \leq \mathbb{E}[M_{(x,y),\nu}] = \pi_{(i)}\delta + \pi_{(ii)}\nu_{(ii)}(x, y)$ and thus $\nu_{(ii)}(x, y) \geq \delta$ (and of course, $\nu_{(i)}(x, y) = 0$). If $F(x, y) = 1$, we have $\delta \leq \mathbb{E}[M_{(x,y),\nu}] = \pi_{(i)}\nu_{(i)}(x, y) + \pi_{(ii)}(1 - \nu_{(ii)}(x, y))\frac{\delta}{1-\delta/4}$ and thus either $\nu_{(i)}(x, y) \geq \delta$ or $\nu_{(ii)}(x, y) \leq \delta/4$. For $k = O(\log(1/\delta) + \log n)$, if we sample rectangles $\{R^w : w \in \{0, 1\}^k\}, \{Q^s : s \in \{0, 1\}^k\}$ where each $R^w \sim \nu_{(i)}$ independently and each $Q^s \sim \nu_{(ii)}$ independently, then with positive probability we arrive at a pair of multisets such that if $F(x, y) = 0$ then $(x, y) \notin \bigcup_w R^w$ and (x, y) is in $\geq 2\delta/3$ fraction of Q^s 's, and if $F(x, y) = 1$ then either $(x, y) \in \bigcup_w R^w$ or (x, y) is in $\leq \delta/3$ fraction of Q^s 's. This constitutes an $(\text{NP} \odot \text{coSBP})$ -type protocol for F , with $\alpha := 2\delta/3$.

If $\pi_{(ii)} = 0$ then the above argument yields an NP-type protocol for F (as $\nu_{(i)}(x, y)$ is 0 if $F(x, y) = 0$, and is $\geq \delta$ if $F(x, y) = 1$, so we get that $\bigcup_w R^w$ covers none of $F^{-1}(0)$ and all of $F^{-1}(1)$); in particular, this yields an $(\text{NP} \odot \text{coSBP})$ -type protocol (by having, say, each of the Q^s 's cover all the inputs). Similarly, if $\pi_{(i)} = 0$ then the above argument yields a coSBP-type protocol for F ; the NP part is not needed. \square

By our characterization, the following inclusions are implicit in [GPW18b]. For completeness, we now provide ‘‘primal’’ arguments for them. We consider the time complexity versions, but the same arguments work for query and communication complexities.

Lemma 7. $\text{ZPP}^{\text{NP}[1]} \subseteq \text{NP} \odot \text{coSBP}$.

Proof. It is slightly more convenient to show $\text{CautiousBPP}^{\text{NP}[1]} \subseteq \text{NP} \odot \text{coSBP}$, so consider a $\text{CautiousBPP}^{\text{NP}[1]}$ algorithm for L , where M uses coin tosses $s \in \{0, 1\}^r$ and each NP oracle query has possible witnesses $w \in \{0, 1\}^k$. Consider an NP-type algorithm N that accepts x iff there exist s and w such that $M_s(x)$ produces q and out with $out(1) = 1$ and the NP oracle verifier accepts (q, w) . Note that if $L(x) = 0$ then $N(x)$ must reject (since M_s cannot err if it makes an oracle query). Consider a coSBP-type algorithm A that samples uniformly random s, w , and bit b , and accepts x iff either

- $b = 0, w = 0^k$, and $M_s(x)$ either directly outputs 0 or produces q and out with $out(0) = 0$, or
- $b = 1, M_s(x)$ produces q and out with $out(1) = 0$, and the NP oracle verifier accepts (q, w) .

If $L(x) = 0$ then for every s such that $M_s(x)$ outputs 0, $A(x)$ accepts with probability $\geq 2^{-k}/2$ (by sampling $b = 0$ and $w = 0^k$ if $M_s(x)$ outputs 0 either directly or after the oracle responds 0, or by sampling $b = 1$ and a correct witness w if $M_s(x)$ outputs 0 after the oracle responds 1); thus overall $A(x)$ accepts with probability $\geq \alpha := 2^{-k} \cdot 3/8$. If $L(x) = 1$ and $N(x)$ rejects, then $A(x)$ can only accept if $b = 0, w = 0^k$, and $M_s(x)$ directly outputs 0, so $A(x)$ accepts with probability $\leq 2^{-k}/8 \leq \alpha/2$. Thus N and A together show that $L \in \text{NP} \odot \text{coSBP}$. \square

Lemma 8. $\text{NP} \odot \text{coSBP} \subseteq \text{PostBPP}$.

Proof. A PostBPP-type algorithm is randomized, produces an output from $\{0, 1, \perp\}$, and must be correct with probability $\geq 3/4$ conditioned on not outputting \perp . [OS18] showed that $\text{PostBPP} = \text{P}_{\parallel}^{\text{SBP}}$, and the proof works for partial functions (in which case, if the algorithm makes an oracle

query outside the SBP function’s domain, the algorithm is required to be correct for both possible responses from the oracle). We have $\text{NP} \oplus \text{coSBP} \subseteq \text{P}_{\parallel}^{\text{SBP}[2]}$ by using one oracle query to evaluate the NP function and another to evaluate the complement of the coSBP function, and outputting 1 iff the first query returns 1 or the second query returns 0. \square

Since $\text{ZPP}^{\text{NP}[1]}$ and PostBPP are closed under complement, Lemma 7 and Lemma 8 imply that $\text{ZPP}^{\text{NP}[1]} \subseteq \text{coNP} \oplus \text{SBP} \subseteq \text{PostBPP}$, as shown in Figure 2.

Acknowledgments

I thank Mika Göös and Toniann Pitassi for discussions, and anonymous reviewers for thorough comments and suggestions.

References

- [AA18] Scott Aaronson and Andris Ambainis. Forrelation: A problem that optimally separates quantum from classical computing. *SIAM Journal on Computing*, 47(3):982–1038, 2018. doi:10.1137/15M1050902.
- [ABB⁺16] Anurag Anshu, Aleksandrs Belovs, Shalev Ben-David, Mika Göös, Rahul Jain, Robin Kothari, Troy Lee, and Miklos Santha. Separations in communication complexity using cheat sheets and information complexity. In *Proceedings of the 57th Symposium on Foundations of Computer Science (FOCS)*, pages 555–564. IEEE, 2016. doi:10.1109/FOCS.2016.66.
- [ABK16] Scott Aaronson, Shalev Ben-David, and Robin Kothari. Separations in query complexity using cheat sheets. In *Proceedings of the 48th Symposium on Theory of Computing (STOC)*, pages 863–876. ACM, 2016. doi:10.1145/2897518.2897644.
- [AKK16] Andris Ambainis, Martins Kokainis, and Robin Kothari. Nearly optimal separations between communication (or query) complexity and partitions. In *Proceedings of the 31st Computational Complexity Conference (CCC)*, pages 4:1–4:14. Schloss Dagstuhl, 2016. doi:10.4230/LIPIcs.CCC.2016.4.
- [BCH⁺17] Adam Bouland, Lijie Chen, Dhiraaj Holden, Justin Thaler, and Prashant Vasudevan. On the power of statistical zero knowledge. In *Proceedings of the 58th Symposium on Foundations of Computer Science (FOCS)*, pages 708–719. IEEE, 2017. doi:10.1109/FOCS.2017.71.
- [BDHT17] Shalev Ben-David, Pooya Hatami, and Avishay Tal. Low-sensitivity functions from unambiguous certificates. In *Proceedings of the 8th Innovations in Theoretical Computer Science Conference (ITCS)*, pages 28:1–28:23. Schloss Dagstuhl, 2017. doi:10.4230/LIPIcs.ITCS.2017.28.
- [BdW02] Harry Buhrman and Ronald de Wolf. Complexity measures and decision tree complexity: A survey. *Theoretical Computer Science*, 288(1):21–43, 2002. doi:10.1016/S0304-3975(01)00144-X.

- [BEGJ00] Maria Luisa Bonet, Juan Luis Esteban, Nicola Galesi, and Jan Johannsen. On the relative complexity of resolution refinements and cutting planes proof systems. *SIAM Journal on Computing*, 30(5):1462–1484, 2000. doi:10.1137/S0097539799352474.
- [BR17] Yakov Babichenko and Aviad Rubinfeld. 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.
- [BVd07] Harry Buhrman, Nikolai Vereshchagin, and Ronald de Wolf. On computation and communication with small bias. In *Proceedings of the 22nd Conference on Computational Complexity (CCC)*, pages 24–32. IEEE, 2007. doi:10.1109/CCC.2007.18.
- [CC06] Jin-Yi Cai and Venkatesan Chakaravathy. On zero error algorithms having oracle access to one query. *Journal of Combinatorial Optimization*, 11(2):189–202, 2006. doi:10.1007/s10878-006-7130-0.
- [CKLM18] Arkadev Chattopadhyay, Michal Koucký, Bruno Loff, and Sagnik Mukhopadhyay. Simulation beats richness: New data-structure lower bounds. In *Proceedings of the 50th Symposium on Theory of Computing (STOC)*, pages 1013–1020. ACM, 2018. doi:10.1145/3188745.3188874.
- [CKLM19] Arkadev Chattopadhyay, Michal Koucký, Bruno Loff, and Sagnik Mukhopadhyay. Simulation theorems via pseudo-random properties. *Computational Complexity*, 28(4):617–659, 2019. doi:10.1007/s00037-019-00190-7.
- [CLRS16] Siu On Chan, James Lee, Prasad Raghavendra, and David Steurer. Approximate constraint satisfaction requires large LP relaxations. *Journal of the ACM*, 63(4):34:1–34:22, 2016. doi:10.1145/2811255.
- [CP08] Richard Chang and Suresh Purini. Amplifying ZPP^{SAT[1]} and the two queries problem. In *Proceedings of the 23rd Conference on Computational Complexity (CCC)*, pages 41–52. IEEE, 2008. doi:10.1109/CCC.2008.32.
- [dNV16] 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 57th Symposium on Foundations of Computer Science (FOCS)*, pages 295–304. IEEE, 2016. doi:10.1109/FOCS.2016.40.
- [GGKS18] 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.
- [GJPW18] Mika Göös, T.S. Jayram, Toniann Pitassi, and Thomas Watson. Randomized communication vs. partition number. *ACM Transactions on Computation Theory*, 10(1):4:1–4:20, 2018. doi:10.1145/3170711.
- [GKPW19] Mika Göös, Pritish Kamath, Toniann Pitassi, and Thomas Watson. Query-to-communication lifting for P^{NP}. *Computational Complexity*, 28(1):113–144, 2019. doi:10.1007/s00037-018-0175-5.

- [GKRS19] Mika Göös, Prithish Kamath, Robert Robere, and Dmitry Sokolov. Adventures in monotone complexity and TFNP. In *Proceedings of the 10th Innovations in Theoretical Computer Science Conference (ITCS)*, pages 38:1–38:19. Schloss Dagstuhl, 2019. doi:10.4230/LIPIcs.ITCS.2019.38.
- [GLM⁺16] Mika Göös, Shachar Lovett, Raghu Meka, Thomas Watson, and David Zuckerman. Rectangles are nonnegative juntas. *SIAM Journal on Computing*, 45(5):1835–1869, 2016. doi:10.1137/15M103145X.
- [Göö15] Mika Göös. Lower bounds for clique vs. independent set. In *Proceedings of the 56th Symposium on Foundations of Computer Science (FOCS)*, pages 1066–1076. IEEE, 2015. doi:10.1109/FOCS.2015.69.
- [GP18] Mika Göös and Toniann Pitassi. Communication lower bounds via critical block sensitivity. *SIAM Journal on Computing*, 47(5):1778–1806, 2018. doi:10.1137/16M1082007.
- [GPW17] Mika Göös, Toniann Pitassi, and Thomas Watson. Query-to-communication lifting for BPP. In *Proceedings of the 58th Symposium on Foundations of Computer Science (FOCS)*, pages 132–143. IEEE, 2017. doi:10.1109/FOCS.2017.21.
- [GPW18a] Mika Göös, Toniann Pitassi, and Thomas Watson. Deterministic communication vs. partition number. *SIAM Journal on Computing*, 47(6):2435–2450, 2018. Special issue for selected papers from FOCS 2015. doi:10.1137/16M1059369.
- [GPW18b] Mika Göös, Toniann Pitassi, and Thomas Watson. The landscape of communication complexity classes. *Computational Complexity*, 27(2):245–304, 2018. doi:10.1007/s00037-018-0166-6.
- [GR18] Mika Göös and Aviad Rubinfeld. Near-optimal communication lower bounds for approximate Nash equilibria. In *Proceedings of the 59th Symposium on Foundations of Computer Science (FOCS)*, pages 397–403. IEEE, 2018. doi:10.1109/FOCS.2018.00045.
- [GW16] Mika Göös and Thomas Watson. Communication complexity of set-disjointness for all probabilities. *Theory of Computing*, 12(1):1–23, 2016. Special issue for selected papers from APPROX–RANDOM 2014. doi:10.4086/toc.2016.v012a009.
- [HR90] 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)90027-1.
- [IW10] Russell Impagliazzo and Ryan Williams. Communication complexity with synchronized clocks. In *Proceedings of the 25th Conference on Computational Complexity (CCC)*, pages 259–269. IEEE, 2010. doi:10.1109/CCC.2010.32.
- [Joh01] Jan Johannsen. Depth lower bounds for monotone semi-unbounded fan-in circuits. *RAIRO - Theoretical Informatics and Applications*, 35:277–286, 2001. doi:10.1051/ita:2001120.
- [Juk12] Stasys Jukna. *Boolean Function Complexity: Advances and Frontiers*, volume 27 of *Algorithms and Combinatorics*. Springer, 2012.

- [Kla03] Hartmut Klauck. Rectangle size bounds and threshold covers in communication complexity. In *Proceedings of the 18th Conference on Computational Complexity (CCC)*, pages 118–134. IEEE, 2003. doi:10.1109/CCC.2003.1214415.
- [Kla07] Hartmut Klauck. Lower bounds for quantum communication complexity. *SIAM Journal on Computing*, 37(1):20–46, 2007. doi:10.1137/S0097539702405620.
- [Kla11] Hartmut Klauck. On Arthur Merlin games in communication complexity. In *Proceedings of the 26th Conference on Computational Complexity (CCC)*, pages 189–199. IEEE, 2011. doi:10.1109/CCC.2011.33.
- [KMR17] Pravesh Kothari, Raghu Meka, and Prasad Raghavendra. Approximating rectangles by juntas and weakly-exponential lower bounds for LP relaxations of CSPs. In *Proceedings of the 49th Symposium on Theory of Computing (STOC)*, pages 590–603. ACM, 2017. doi:10.1145/3055399.3055438.
- [KN97] Eyal Kushilevitz and Noam Nisan. *Communication Complexity*. Cambridge University Press, 1997.
- [LRS15] James Lee, Prasad Raghavendra, and David Steurer. Lower bounds on the size of semidefinite programming relaxations. In *Proceedings of the 47th Symposium on Theory of Computing (STOC)*, pages 567–576. ACM, 2015. doi:10.1145/2746539.2746599.
- [New91] Ilan Newman. Private vs. common random bits in communication complexity. *Information Processing Letters*, 39(2):67–71, 1991. doi:10.1016/0020-0190(91)90157-D.
- [OS18] Ryan O’Donnell and A. C. Cem Say. The weakness of CTC qubits and the power of approximate counting. *ACM Transactions on Computation Theory*, 10(2):5:1–5:22, 2018. doi:10.1145/3196832.
- [PR17] 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.
- [PR18] 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.
- [PSS14] Periklis Papakonstantinou, Dominik Scheder, and Hao Song. Overlays and limited memory communication. In *Proceedings of the 29th Conference on Computational Complexity (CCC)*, pages 298–308. IEEE, 2014. doi:10.1109/CCC.2014.37.
- [PSW20] Toniann Pitassi, Morgan Shirley, and Thomas Watson. Nondeterministic and randomized boolean hierarchies in communication complexity. In *Proceedings of the 47th International Colloquium on Automata, Languages, and Programming (ICALP)*, pages 92:1–92:19. Schloss Dagstuhl, 2020. doi:10.4230/LIPIcs.ICALP.2020.92.
- [RM99] Ran Raz and Pierre McKenzie. Separation of the monotone NC hierarchy. *Combinatorica*, 19(3):403–435, 1999. doi:10.1007/s004930050062.

- [RPRC16] 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.
- [RW16] 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.
- [RY20] Anup Rao and Amir Yehudayoff. *Communication Complexity and Applications*. Cambridge University Press, 2020.
- [She11] Alexander Sherstov. The pattern matrix method. *SIAM Journal on Computing*, 40(6):1969–2000, 2011. doi:10.1137/080733644.
- [Tri10] Rahul Tripathi. The 1-versus-2 queries problem revisited. *Theory of Computing Systems*, 46(2):193–221, 2010. doi:10.1007/s00224-008-9126-x.
- [Ver99] Nikolai Vereshchagin. Relativizability in complexity theory. In *Provability, Complexity, Grammars*, volume 192 of *AMS Translations, Series 2*, pages 87–172. American Mathematical Society, 1999.
- [Wat19] Thomas Watson. Amplification with one NP oracle query. In *Proceedings of the 46th International Colloquium on Automata, Languages, and Programming (ICALP), Track A*, pages 96:1–96:13. Schloss Dagstuhl, 2019. doi:10.4230/LIPIcs.ICALP.2019.96.
- [WYY17] Xiaodi Wu, Penghui Yao, and Henry Yuen. Raz–McKenzie simulation with the inner product gadget. Technical Report TR17-010, Electronic Colloquium on Computational Complexity (ECCC), 2017. URL: <https://eccc.weizmann.ac.il/report/2017/010/>.