



Nondeterministic and Randomized Boolean Hierarchies in Communication Complexity

Toniann Pitassi
*University of Toronto and
 Institute for Advanced Study*

Morgan Shirley
University of Toronto

Thomas Watson
University of Memphis

March 8, 2019

Abstract

We study the Boolean Hierarchy in the context of two-party communication complexity, as well as the analogous hierarchy defined with one-sided error randomness instead of nondeterminism. Our results provide a complete picture of the relationships among complexity classes within and across these two hierarchies. In particular, we prove a query-to-communication lifting theorem for all levels of the Boolean Hierarchy and use it to resolve an open problem stated in the paper by Halstenberg and Reischuk (CCC 1988) which initiated this topic.

1 Introduction

The Boolean Hierarchy in classical complexity theory consists of problems that have a polynomial-time algorithm making a constant number of queries to an NP oracle. This hierarchy has an intricate relationship with other complexity classes, and its second level (DP) captures the complexity of certain “exact” versions of optimization problems. It consists of an infinite sequence of complexity classes $\text{NP}(q)$ for $q = 1, 2, 3, \dots$ (where $\text{NP}(1) = \text{NP}$ and $\text{NP}(2) = \text{DP}$). Among the several equivalent ways of defining these classes [Wec85, CH86, KSW87, Wag88], perhaps the simplest is that $\text{NP}(q)$ consists of all decision problems that can be computed by taking the parity of the answers to a sequence of q NP problems on the given input. As illustrated in Figure 1, it is known that these levels are intertwined with the classes $\text{P}_{\parallel}^{\text{NP}[q]}$ of all decision problems solvable in polynomial time using q nonadaptive NP queries (for constant q) [KSW87, Wag88, Bei91]:

$$\text{NP}(q) \subseteq \text{P}_{\parallel}^{\text{NP}[q]} \subseteq \text{NP}(q+1) \quad \text{and} \quad \text{coNP}(q) \subseteq \text{P}_{\parallel}^{\text{NP}[q]} \subseteq \text{coNP}(q+1)$$

(by closure of $\text{P}_{\parallel}^{\text{NP}[q]}$ under complementation). Here, $\text{coNP}(q)$ means $\text{co}(\text{NP}(q))$ rather than $(\text{coNP})(q)$.

Analogous to the above *Nondeterministic* Boolean Hierarchy, one can define the *Randomized* Boolean Hierarchy by using RP (one-sided error randomized polynomial time) instead of NP in the definitions [BBJ⁺89]. The analogous inclusions like in Figure 1 hold among all the classes $\text{RP}(q)$, $\text{coRP}(q)$, and $\text{P}_{\parallel}^{\text{RP}[q]}$, by similar arguments. Although the (suitably defined) *Polynomial* Hierarchy over RP is known to collapse to its second level, which equals BPP [ZH86], the *Boolean* Hierarchy over RP has not been widely studied.

We investigate the Nondeterministic and Randomized Boolean Hierarchies in the context of two-party communication complexity. In the basic deterministic model of communication [Yao79, KN97], Alice is given an input x and Bob is given an input y , and they wish to collaboratively evaluate some function $F(x, y)$ of their joint input by engaging in a protocol that specifies how they exchange bits of information about their inputs. Many classical complexity classes (P, RP, NP, and so on) have two-party communication analogues [BFS86]. The area of structural communication complexity, which concerns the properties of and relationships among these classes, is undergoing a renaissance and has turned out to yield new techniques and perspectives for understanding questions in a variety of other areas (circuit complexity, proof complexity, data structures, learning theory, delegation, fine-grained complexity, property testing, cryptography, extended formulations, etc.) [GPW18b]. For any classical time-bounded complexity class \mathcal{C} , we use \mathcal{C}^{cc} to denote its

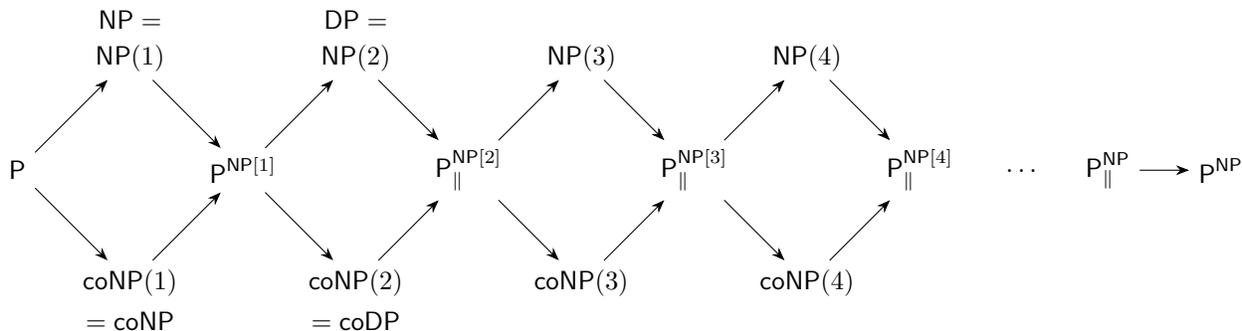


Figure 1: Relations between classes in the Boolean Hierarchy. Here, $\mathcal{C}_1 \rightarrow \mathcal{C}_2$ represents $\mathcal{C}_1 \subseteq \mathcal{C}_2$.

communication complexity analogue—the class of all two-party functions on n bits that admit a protocol communicating at most $\text{polylog}(n)$ bits, in a model defined by analogy with the classical \mathcal{C} .

Halstenberg and Reischuk [HR88, HR90] initiated the study of the Nondeterministic Boolean Hierarchy in two-party communication complexity. They observed that the inclusions shown in Figure 1 hold for the communication versions of the classes, by essentially the same proofs as in the time-bounded setting. They also proved that $\text{NP}(q)^{\text{cc}} \neq \text{coNP}(q)^{\text{cc}}$, which simultaneously implies that each of the inclusions is strict: $\text{NP}(q)^{\text{cc}} \subsetneq P_{\parallel}^{\text{NP}[q]^{\text{cc}}} \subsetneq \text{NP}(q+1)^{\text{cc}}$.

The communication version of the Randomized Boolean Hierarchy has not been explicitly studied as far as we know, but it is interesting since $\text{EQUALITY} \in \text{coRP}^{\text{cc}}$ and many randomized protocols have been designed by reduction to this fact (e.g., $\text{GREATERTHAN} \in P^{\text{EQUALITY}} \subseteq P^{\text{RP}^{\text{cc}}}$). Recently, it was shown that EQUALITY is not the “only” total two-party function randomness is good for: $\text{RP}^{\text{cc}} \not\subseteq P^{\text{EQUALITY}}$ [CLV18]. What can we say about the power of a fixed number of queries to an RP^{cc} oracle? Our first contribution strengthens the aforementioned separation due to Halstenberg and Reischuk.

Theorem 1. *For total functions, $\text{coRP}(q)^{\text{cc}} \not\subseteq \text{NP}(q)^{\text{cc}}$ for every constant q .*

Since $\text{RP}^{\text{cc}} \subseteq \text{NP}^{\text{cc}}$, Theorem 1 simultaneously implies that each of the inclusions in the Randomized Boolean Hierarchy is strict: $\text{RP}(q)^{\text{cc}} \subsetneq P_{\parallel}^{\text{RP}[q]^{\text{cc}}} \subsetneq \text{RP}(q+1)^{\text{cc}}$, and thus the hierarchy does not collapse. Previously, no separation beyond the first level seemed to be known in the literature. Our proof of Theorem 1 is completely different from (and more involved than) Halstenberg and Reischuk’s proof of $\text{coNP}(q)^{\text{cc}} \not\subseteq \text{NP}(q)^{\text{cc}}$, which used the “easy-hard argument” of [Kad88].

In [HR88, HR90], Halstenberg and Reischuk also explicitly asked whether the inclusion $P_{\parallel}^{\text{NP}[q]^{\text{cc}}} \subseteq \text{NP}(q+1)^{\text{cc}} \cap \text{coNP}(q+1)^{\text{cc}}$ is strict. When $q = 0$, this is answered by the familiar results that $P^{\text{cc}} = \text{NP}^{\text{cc}} \cap \text{coNP}^{\text{cc}}$ when the classes are defined to contain only total functions [HR93], whereas $P^{\text{cc}} \subsetneq \text{NP}^{\text{cc}} \cap \text{coNP}^{\text{cc}}$ (indeed, $P^{\text{cc}} \subsetneq \text{ZPP}^{\text{cc}}$) holds when partial functions (promise problems) are allowed. For $q > 0$, we resolve this 31-year-old open question by proving that the situation is analogous to the $q = 0$ case.

Theorem 2. *For total functions,*

$$P_{\parallel}^{\text{NP}[q]^{\text{cc}}} = \text{NP}(q+1)^{\text{cc}} \cap \text{coNP}(q+1)^{\text{cc}} \quad \text{and} \quad P_{\parallel}^{\text{RP}[q]^{\text{cc}}} = \text{RP}(q+1)^{\text{cc}} \cap \text{coRP}(q+1)^{\text{cc}}$$

for every constant q .

Theorem 3. *For partial functions, $\text{RP}(q+1)^{\text{cc}} \cap \text{coRP}(q+1)^{\text{cc}} \not\subseteq P_{\parallel}^{\text{NP}[q]^{\text{cc}}}$ for every constant q .*

Since $\text{RP}^{\text{cc}} \subseteq \text{NP}^{\text{cc}}$, Theorem 3 implies that

$$P_{\parallel}^{\text{NP}[q]^{\text{cc}}} \subsetneq \text{NP}(q+1)^{\text{cc}} \cap \text{coNP}(q+1)^{\text{cc}} \quad \text{and} \quad P_{\parallel}^{\text{RP}[q]^{\text{cc}}} \subsetneq \text{RP}(q+1)^{\text{cc}} \cap \text{coRP}(q+1)^{\text{cc}}$$

for partial functions. Taken together, Theorem 1, Theorem 2, and Theorem 3 form a complete picture of the relationships among the classes within and across the Nondeterministic and Randomized Boolean Hierarchies in communication complexity, for both total and partial functions.

Our proof of [Theorem 3](#) uses the paradigm of *query-to-communication lifting* [[RM99](#), [GLM⁺16](#), [Göös15](#), [GPW18a](#), [GKPW17](#), [GPW17](#), [Wat19](#)]. This approach to proving communication lower bounds has led to breakthroughs on fundamental questions in communication complexity and many of its application areas. The idea consists of two steps:

- (1) First prove an analogous lower bound in the simpler setting of decision tree depth complexity (a.k.a. query complexity). This step captures the combinatorial core of the lower bound argument without the burden of dealing with the full power of communication protocols.
- (2) Then apply a *lifting theorem*, which translates the query lower bound into a communication lower bound for a related two-party function. This step encapsulates the general-purpose machinery for dealing with protocols, and can be reused from one argument to the next.

The availability of a lifting theorem greatly simplifies the task of proving certain communication lower bounds, because it divorces the problem-specific aspects from the generic aspects. The format of a lifting theorem is that if $f: \{0, 1\}^n \rightarrow \{0, 1\}$ is any partial function and $g: \mathcal{X} \times \mathcal{Y} \rightarrow \{0, 1\}$ is a certain “small” two-party function called a gadget, then the communication complexity of the two-party composed function $f \circ g^n: \mathcal{X}^n \times \mathcal{Y}^n \rightarrow \{0, 1\}$ —in which Alice gets $x = (x_1, \dots, x_n)$, Bob gets $y = (y_1, \dots, y_n)$, and their goal is to evaluate $(f \circ g^n)(x, y) := f(g(x_1, y_1), \dots, g(x_n, y_n))$ —should be approximately the query complexity of the outer function f . One direction is generally straightforward: given a query upper bound for f , a communication upper bound for $f \circ g^n$ is witnessed by a protocol that simulates the decision tree for f and evaluates $g(x_i, y_i)$ whenever it queries the i^{th} bit of the input to f ; the number of bits of communication is at most the number of queries made by the decision tree times the (small) cost of evaluating a copy of g . The other direction is the challenging part: despite Alice and Bob’s ability to send messages that depend in arbitrary ways on all n coordinates, they nevertheless cannot do much better than just simulating a decision tree, which involves “talking about” one coordinate at a time.

A lifting theorem must be stated with respect to a particular model of computation, such as deterministic, one-sided error randomized, nondeterministic, etc., which we associate with the corresponding complexity classes. Indeed, lifting theorems are known for P [[RM99](#), [GPW18a](#)], RP [[GPW17](#)], NP [[GLM⁺16](#), [Göös15](#)], and many other classes. It is convenient to recycle complexity class names to denote the complexity of a given function in the corresponding model, e.g., $P^{\text{dt}}(f)$ is the minimum worst-case number of queries made by any decision tree that computes f , and $P^{\text{cc}}(F)$ is the minimum worst-case communication cost of any protocol that computes F . With this notation, the deterministic lifting theorem from [[RM99](#), [GPW18a](#)] can be stated as: for all f , $P^{\text{cc}}(f \circ g^n) = P^{\text{dt}}(f) \cdot \Theta(\log n)$ where $g: [m] \times \{0, 1\}^m \rightarrow \{0, 1\}$ is the “index” gadget defined by $g(x, y) = y_x$ with $m := n^{20}$. (Note that $P^{\text{cc}}(g) = O(\log n)$ since Alice can send her $\log m$ -bit “pointer” to Bob, who responds with the pointed-to bit from his string.) The index gadget has also been used in lifting theorems for several other complexity classes.

We prove lifting theorems for all classes in the Nondeterministic Boolean Hierarchy, with the index gadget.

Theorem 4. *For every partial function $f: \{0, 1\}^n \rightarrow \{0, 1\}$ and every constant q ,*

$$(i) \text{ NP}(q)^{\text{cc}}(f \circ g^n) = \text{NP}(q)^{\text{dt}}(f) \cdot \Theta(\log n)$$

$$(ii) P_{\parallel}^{\text{NP}[q]\text{cc}}(f \circ g^n) = P_{\parallel}^{\text{NP}[q]\text{dt}}(f) \cdot \Theta(\log n)$$

where $g: [m] \times \{0, 1\}^m \rightarrow \{0, 1\}$ is the index gadget defined by $g(x, y) = y_x$ with $m := n^{20}$.

Only part (ii) is needed for proving [Theorem 3](#), but part (i) forms an ingredient in the proof of (ii) and is of independent interest.

The most closely related lifting theorem to [Theorem 4](#) is the one for P^{NP} [[GKPW17](#)], corresponding to computations that make an unbounded number of adaptive queries to an NP oracle. In that paper, the overall idea was to approximately characterize P^{NP} complexity in terms of *decision lists* (DL), and then prove a lifting theorem directly for DLs. Briefly, a conjunction DL (introduced by [[Riv87](#)]) is a sequence of small-width conjunctions each with an associated output bit, and the output is determined by finding the first conjunction in the list that accepts the given input. A rectangle DL is similar but with combinatorial rectangles instead of conjunctions. The proof from [[GKPW17](#)] shows how to convert a rectangle DL for $f \circ g^n$ into a conjunction DL for f .

The gist of our arguments for both parts of [Theorem 4](#) is to approximately characterize (via different techniques than for P^{NP}) these classes in terms of DLs with a bounded number of *alternations* (how many times the associated output bit flips as we walk down the entire DL). The DL lifting argument from [\[GKPW17\]](#) does not preserve the number of alternations, but we show how it can be adapted to do this. Our techniques also yield an approximate lifting theorem for $\mathsf{P}_{\parallel}^{\mathsf{NP}}$ (corresponding to computations that make an unbounded number of nonadaptive NP oracle queries), but we omit the details.

2 Preliminaries

We assume familiarity with deterministic computation in query and communication complexity [\[Juk12, KN97\]](#). Recall the following standard definitions of nondeterministic and one-sided error randomized models:

- An NP^{dt} decision tree is a DNF formula Φ . Given an input z , the output of such a decision tree is Φ evaluated on z . A function f is computed by Φ if $f(z) = \Phi(z)$ on all inputs z for which $f(z)$ is defined. The cost of Φ is the maximum width (number of literals) in any conjunction in Φ .
- An NP^{cc} protocol is a set \mathcal{R} of combinatorial rectangles. Given an input (x, y) , the output of such a protocol is $\mathcal{R}(x, y) := 1$ iff there exists an $R \in \mathcal{R}$ containing (x, y) . A two-party function F is computed by \mathcal{R} if $F(x, y) = \mathcal{R}(x, y)$ on all inputs (x, y) for which $F(x, y)$ is defined. The cost of \mathcal{R} is $\lceil \log(|\mathcal{R}| + 1) \rceil$, which intuitively represents the number of bits required to specify a rectangle in \mathcal{R} or indicate that the input is in no such rectangle.
- An RP^{dt} decision tree is a distribution \mathbf{T} over deterministic decision trees. Given an input z , the output of such a decision tree is computed by sampling a deterministic decision tree T from \mathbf{T} and evaluating $T(z)$. A function f is computed by \mathbf{T} if for all $z \in f^{-1}(0)$, $\Pr_{T \sim \mathbf{T}}[T(z) = 1] = 0$ and for all $z \in f^{-1}(1)$, $\Pr_{T \sim \mathbf{T}}[T(z) = 1] \geq 1/2$. The cost of \mathbf{T} is the maximum number of bits queried in any T in its support.
- An RP^{cc} protocol is a distribution $\mathbf{\Pi}$ over deterministic communication protocols. Given an input (x, y) , the output of such a protocol is computed by sampling a deterministic protocol Π from $\mathbf{\Pi}$ and evaluating $\Pi(x, y)$. A function F is computed by $\mathbf{\Pi}$ if for all $(x, y) \in F^{-1}(0)$, $\Pr_{\Pi \sim \mathbf{\Pi}}[\Pi(x, y) = 1] = 0$ and for all $(x, y) \in F^{-1}(1)$, $\Pr_{\Pi \sim \mathbf{\Pi}}[\Pi(x, y) = 1] \geq 1/2$. The cost of $\mathbf{\Pi}$ is the maximum number of bits exchanged in any Π in its support.

Let \mathcal{C} be an arbitrary complexity class name representing a model of computation (such as NP or RP). We let $\mathcal{C}^{\text{cc}}(F)$ denote the communication complexity of a two-party function F in the corresponding model: the minimum cost of any \mathcal{C}^{cc} protocol computing F . We let $\mathcal{C}^{\text{dt}}(f)$ denote the query complexity of a Boolean function f in the corresponding model: the minimum cost of any \mathcal{C}^{dt} decision tree computing f . Often we will abuse notation by having F or f refer to an infinite *family* of functions, where there is at most one function in the family for each possible input length. In this case, $\mathcal{C}^{\text{cc}}(F)$ or $\mathcal{C}^{\text{dt}}(f)$ will be the complexity parameterized by the input length n ; we typically express this with asymptotic notation. When written by itself, \mathcal{C}^{cc} or \mathcal{C}^{dt} denotes the class of all families of functions with complexity at most polylogarithmic in n , in the corresponding model. We will always clarify whether a class \mathcal{C}^{cc} or \mathcal{C}^{dt} is meant to contain partial functions or just total functions, since this is not explicit in the notation.

For RP^{cc} and RP^{dt} , the constant $1/2$ in the success probability is arbitrary: by amplification, choosing a different positive constant in the definition would only affect the complexity of any function by a constant factor. Also note that $\mathsf{NP}^{\text{dt}}(f) \leq \mathsf{RP}^{\text{dt}}(f)$ for all f , and since we defined RP^{cc} using the public-coin model, we have $\mathsf{NP}^{\text{cc}}(F) \leq \mathsf{RP}^{\text{cc}}(F) + O(\log n)$ for all F (by decreasing the number of random bits for sampling a protocol to $O(\log n)$ and using nondeterminism to guess an outcome that results in output 1).

2.1 Nondeterministic and Randomized Boolean Hierarchies

In the following definitions, restrict \mathcal{C} to be either NP or RP. We will use two different but equivalent definitions of the constituent levels of the Nondeterministic and Randomized Boolean Hierarchies. Our “official” definition is in terms of the following “decision list functions”:

Definition 5. $\Delta_q: \{0, 1\}^q \rightarrow \{0, 1\}$ is defined inductively as follows:

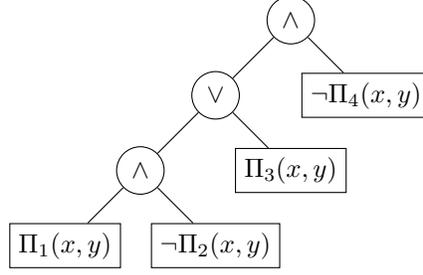


Figure 2: A visualization of a $\mathcal{C}(4)^{\text{cc}}$ protocol, where each Π_i is a \mathcal{C}^{cc} protocol.

- $\Delta_1(z_1) := z_1$.
- If q is odd, $\Delta_q(z_1, \dots, z_{q-1}, z_q) := \Delta_{q-1}(z_1, \dots, z_{q-1}) \vee z_q$.
- If q is even, $\Delta_q(z_1, \dots, z_{q-1}, z_q) := \Delta_{q-1}(z_1, \dots, z_{q-1}) \wedge (\neg z_q)$.

In other words, letting $\oplus: \mathbb{N} \rightarrow \{0, 1\}$ denote the parity function, we have $\Delta_q(z) := \oplus(i)$ where i is the greatest index such that $z_i = 1$ (or $i = 0$ if z is all zeros).

Definition 6. A $\mathcal{C}(q)^{\text{cc}}$ protocol is an ordered list of q many \mathcal{C}^{cc} protocols $\Pi = (\Pi_1, \dots, \Pi_q)$. Given an input (x, y) , the output of the protocol is $\Pi(x, y) := \Delta_q(\Pi_1(x, y), \dots, \Pi_q(x, y))$. The cost of a $\mathcal{C}(q)^{\text{cc}}$ protocol is the sum of the costs of the component \mathcal{C}^{cc} protocols.

See Figure 2 for a visualization. The Nondeterministic Boolean Hierarchy is $\bigcup_{\text{constant } q} \text{NP}(q)^{\text{cc}}$ and the Randomized Boolean Hierarchy is $\bigcup_{\text{constant } q} \text{RP}(q)^{\text{cc}}$. We are also interested in the complement classes $\text{coNP}(q)^{\text{cc}}$ and $\text{coRP}(q)^{\text{cc}}$. As is standard, when we write $\text{co}\mathcal{C}(q)^{\text{cc}}$ we refer to the class $\text{co}\mathcal{C}(q)^{\text{cc}}$ (that is, functions that are the negations of functions in $\mathcal{C}(q)^{\text{cc}}$) as opposed to $(\text{co}\mathcal{C})(q)^{\text{cc}}$ (that is, where the component protocols are $\text{co}\mathcal{C}^{\text{cc}}$ protocols).

There are analogous definitions in query complexity:

Definition 7. A $\mathcal{C}(q)^{\text{dt}}$ decision tree is an ordered list of q many \mathcal{C}^{dt} decision trees $T = (T_1, \dots, T_q)$. Given an input z , the output of the decision tree is $T(z) := \Delta_q(T_1(z), \dots, T_q(z))$. The cost of a $\mathcal{C}(q)^{\text{dt}}$ decision tree is the sum of the costs of the component \mathcal{C}^{dt} decision trees.

Our alternative definition of the Nondeterministic and Randomized Boolean Hierarchies simply replaces Δ_q with the parity function $\oplus_q: \{0, 1\}^q \rightarrow \{0, 1\}$. In Appendix A.1 we provide the standard proof that these official and alternative definitions are equivalent:

Lemma 8. For $\mathcal{C} \in \{\text{NP}, \text{RP}\}$, if the definitions of $\mathcal{C}(q)^{\text{cc}}$ and $\mathcal{C}(q)^{\text{dt}}$ are changed to use \oplus_q in place of Δ_q , it only affects the complexity measures $\mathcal{C}(q)^{\text{cc}}(F)$ and $\mathcal{C}(q)^{\text{dt}}(f)$ by a constant factor (depending on q).

We use both definitions in this paper. We found that the Δ_q definition makes it easier to prove the lifting theorems, and the \oplus_q definition makes it easier to prove concrete upper and lower bounds.

2.2 Decision lists

The reason we call Δ_q a “decision list function” is that it highlights the connection between the Boolean Hierarchy classes and the *decision list* models of computation:

Definition 9. A rectangle decision list \mathcal{L}_R is an ordered list of pairs $(R_1, \ell_1), (R_2, \ell_2), \dots$ where each R_i is a combinatorial rectangle, $\ell_i \in \{0, 1\}$ is a label, and the final rectangle in the list contains all inputs in the domain. For an input (x, y) , the output $\mathcal{L}_R(x, y)$ is ℓ_i where i is the first index for which $(x, y) \in R_i$. The cost of \mathcal{L}_R is the log of the length of the list.

Definition 10. A conjunction decision list \mathcal{L}_C is an ordered list of pairs $(C_1, \ell_1), (C_2, \ell_2), \dots$ where each C_i is a conjunction, $\ell_i \in \{0, 1\}$ is a label, and the final conjunction in the list accepts all inputs in the domain. For an input z , the output $\mathcal{L}_C(z)$ is ℓ_i where i is the first index for which $C_i(z) = 1$. The cost of \mathcal{L}_C is the maximum width of any conjunction in the list.

Note that the restriction on the final rectangle/conjunction is without loss of generality. The complexity measures $\text{DL}^{\text{cc}}(F)$ and $\text{DL}^{\text{dt}}(f)$ are the minimum cost of any rectangle/conjunction decision list computing F or f , and the classes DL^{cc} and DL^{dt} contain those functions with complexity at most $\text{polylog}(n)$.

We now define q -alternating decision lists to have the additional restriction that the sequence of output labels ℓ_1, ℓ_2, \dots only flips between 0 and 1 at most q times, and furthermore the last label is 0. This restriction partitions the list into contiguous *levels* where all labels in the same level are equal; without loss of generality the last level consists only of the final “catch-all” entry. For convenience, in the list entries we replace the labels with the level numbers themselves.

Definition 11. A q -alternating rectangle decision list \mathcal{L}_R is an ordered list of pairs $(R_1, \ell_1), (R_2, \ell_2), \dots$ where each R_i is a combinatorial rectangle, $\ell_i \in \{0, 1, \dots, q\}$ is a level such that $\ell_i \geq \ell_{i+1}$ for all i , and the final rectangle in the list contains all inputs in the domain and is the only rectangle at level 0. For an input (x, y) , the output $\mathcal{L}_R(x, y)$ is $\oplus(\ell_i)$ where i is the first index for which $(x, y) \in R_i$. The cost of \mathcal{L}_R is the log of the length of the list.

Definition 12. A q -alternating conjunction decision list \mathcal{L}_C is an ordered list of pairs $(C_1, \ell_1), (C_2, \ell_2), \dots$ where each C_i is a conjunction, $\ell_i \in \{0, 1, \dots, q\}$ is a level such that $\ell_i \geq \ell_{i+1}$ for all i , and the final conjunction in the list accepts all inputs in the domain and is the only conjunction at level 0. For an input z , the output $\mathcal{L}_C(z)$ is $\oplus(\ell_i)$ where i is the first index for which $C_i(z) = 1$. The cost of \mathcal{L}_C is the maximum width of any conjunction in the list.

The complexity measures $\text{DL}(q)^{\text{cc}}(F)$ and $\text{DL}(q)^{\text{dt}}(f)$ are the minimum cost of any q -alternating rectangle/conjunction decision list computing F or f , and the classes $\text{DL}(q)^{\text{cc}}$ and $\text{DL}(q)^{\text{dt}}$ contain those functions with complexity at most $\text{polylog}(n)$.

It turns out that q -alternating decision lists are equivalent to $\text{NP}(q)$ in both communication and query complexity. This follows almost immediately from the definition of Δ_q . For completeness, we include the argument in [Appendix A.2](#).

Lemma 13. $\text{DL}(q)^{\text{cc}}(F) = \Theta(\text{NP}(q)^{\text{cc}}(F))$ and $\text{DL}(q)^{\text{dt}}(f) = \Theta(\text{NP}(q)^{\text{dt}}(f))$ for every constant q . Thus, $\text{DL}(q)^{\text{cc}} = \text{NP}(q)^{\text{cc}}$ and $\text{DL}(q)^{\text{dt}} = \text{NP}(q)^{\text{dt}}$ for partial functions.

This can be contrasted with the lemma from [\[GKPW17\]](#) stating that $\text{DL}^{\text{cc}} = \text{P}^{\text{NP}^{\text{cc}}}$ and $\text{DL}^{\text{dt}} = \text{P}^{\text{NP}^{\text{dt}}}$ for partial functions.

2.3 Parallel queries

In the following definitions, restrict \mathcal{C} to be either NP or RP .

Definition 14. A $\text{P}_{\parallel}^{\mathcal{C}[q]^{\text{cc}}}$ protocol consists of a deterministic protocol Π_{det} that maps an input (x, y) to two things: a function $\text{out}: \{0, 1\}^q \rightarrow \{0, 1\}$ and an ordered list of q many \mathcal{C}^{cc} protocols (Π_1, \dots, Π_q) . The output is then $\text{out}(\Pi_1(x, y), \dots, \Pi_q(x, y))$. The cost of a $\text{P}_{\parallel}^{\mathcal{C}[q]^{\text{cc}}}$ protocol is the communication cost (depth) of Π_{det} plus the maximum over (x, y) of the sum of the costs of the \mathcal{C}^{cc} protocols produced by $\Pi_{\text{det}}(x, y)$.

Definition 15. A $\text{P}_{\parallel}^{\mathcal{C}[q]^{\text{dt}}}$ decision tree consists of a deterministic decision tree T_{det} that maps an input z to two things: a function $\text{out}: \{0, 1\}^q \rightarrow \{0, 1\}$ and an ordered list of q many \mathcal{C}^{dt} decision trees (T_1, \dots, T_q) . The output is then $\text{out}(T_1(z), \dots, T_q(z))$. The cost of a $\text{P}_{\parallel}^{\mathcal{C}[q]^{\text{dt}}}$ decision tree is the query cost (depth) of T_{det} plus the maximum over z of the sum of the costs of the \mathcal{C}^{dt} decision trees produced by $T_{\text{det}}(z)$.

The following lemma states that at each leaf of Π_{det} or T_{det} , we can replace the q “ \mathcal{C} oracle queries” with one “ $\mathcal{C}(q)$ oracle query” (where some leaves may output the oracle’s answer, while other leaves output the negation of it). This was shown in classical time-bounded complexity using the so-called “mind-change argument” [\[Bei91\]](#), and this argument can be translated directly to communication and query complexity. For example, [\[HR90\]](#) used this method to show that $\text{P}_{\parallel}^{\text{NP}[q]^{\text{cc}}} \subseteq \text{NP}(q+1)^{\text{cc}} \cap \text{coNP}(q+1)^{\text{cc}}$. For completeness, we provide the proof in [Appendix A.3](#). We will only need to use the result for $\mathcal{C} = \text{NP}$.

Lemma 16. For $\mathcal{C} \in \{\text{NP}, \text{RP}\}$, we have $\text{P}_{\parallel}^{\mathcal{C}[q]^{\text{cc}}}(F) = \Theta(\text{P}^{\mathcal{C}(q)[1]^{\text{cc}}}(F))$ and $\text{P}_{\parallel}^{\mathcal{C}[q]^{\text{dt}}}(f) = \Theta(\text{P}^{\mathcal{C}(q)[1]^{\text{dt}}}(f))$ for every constant q .

3 Separations

Restatement of Theorem 1. For total functions, $\text{coRP}(q)^{\text{cc}} \not\subseteq \text{NP}(q)^{\text{cc}}$ for every constant q .

In Section 3.1 we prove a weaker version of Theorem 1 that holds for partial functions, by showing the analogous query complexity separation $\text{coRP}(q)^{\text{dt}} \not\subseteq \text{NP}(q)^{\text{dt}}$ then applying our lifting theorem for $\text{NP}(q)$, Theorem 4.(i). (The query complexity separation is known to be false for total functions: $\text{coRP}(q)^{\text{dt}} \subseteq \text{P}^{\text{RP}^{\text{dt}}} \subseteq \text{BPP}^{\text{dt}} \subseteq \text{P}^{\text{dt}} \subseteq \text{NP}(q)^{\text{dt}}$ [Nis91].) This serves two purposes: it is a warmup for our direct proof of Theorem 1 for total functions in Section 3.2, and it forms a component in our proof of Theorem 3 in Section 3.3.

3.1 Proof of Theorem 1 for partial functions

Fix any constant q . Let $\oplus_q \text{GAPOR}: (\{0,1\}^n)^q \rightarrow \{0,1\}$ be the partial function where the input is divided into q blocks $z = (z_1, \dots, z_q)$ having the promise that each $z_i \in \{0,1\}^n$ is either all zeros or at least half ones (call such an input *valid*), and which is defined by $\oplus_q \text{GAPOR}(z) := 1$ iff an odd number of blocks i are such that z_i is nonzero. Note that $\text{RP}(q)^{\text{dt}}(\oplus_q \text{GAPOR}) = O(1)$ by Lemma 8 and the fact that $\text{RP}^{\text{dt}}(\text{GAPOR}) = 1$. Letting $g: [m] \times \{0,1\}^m \rightarrow \{0,1\}$ be the index gadget with $m := (qn)^{20}$, this implies that $\text{RP}(q)^{\text{cc}}(\oplus_q \text{GAPOR} \circ g^{qn}) = O(\log n)$ (here we have used the “easy direction” of $\text{RP}(q)$ lifting; the “hard direction” remains an open problem) and thus $\oplus_q \text{GAPOR} \circ g^{qn} \in \text{coRP}(q)^{\text{cc}}$. We will now prove that $\text{NP}(q)^{\text{dt}}(\oplus_q \text{GAPOR}) = \Omega(n)$, which by Theorem 4.(i) implies that $\text{NP}(q)^{\text{cc}}(\oplus_q \text{GAPOR} \circ g^{qn}) = \Omega(n \log n)$.

Suppose for contradiction $\oplus_q \text{GAPOR}$ has an $\text{NP}(q)^{\text{dt}}$ decision tree of cost $k \leq n/2$, say $T = (\Phi_1, \dots, \Phi_q)$ where each Φ_i is a DNF. By Lemma 8 we may assume the decision tree outputs 1 iff an odd number of these DNFs accept the input, in other words, $T(z) := \oplus_q(\Phi_1(z), \dots, \Phi_q(z))$.

We will iteratively construct a sequence of partial assignments $\rho^j \in (\{0,1,*\}^n)^q$ for $j = 0, \dots, q$ such that (i) there are at least j many values of i for which Φ_i accepts all inputs that are consistent with ρ^j , and (ii) ρ^j is consistent with a valid input where exactly j blocks are nonzero. We obtain the contradiction when $j = q$: by (ii) some valid input z consistent with ρ^q has z_i nonzero for all i and thus $\oplus_q \text{GAPOR}(z) = 1 - \oplus(q)$, yet by (i) we have $\Phi_i(z) = 1$ for all i and thus $T(z) = \oplus(q)$, contradicting the supposed correctness of T .

We will actually maintain stronger invariants than the above (i) and (ii): For (i), we will actually have for some j values of i —we assume they are $1, \dots, j$ for notational convenience—some individual conjunction C_i of Φ_i accepts all inputs consistent with ρ^j . For (ii), ρ^j will actually have the following form: for some fixed assignment $l^j = (l_1, \dots, l_j) \in (\{0,1\}^n)^j$ such that l_i is at least half ones for all $i \in [j]$ (“left”), and for some partial assignment $r^j \in (\{0,*\}^n)^{q-j}$ (“right”), we have $\rho^j := l^j r^j$. The valid input z^j obtained by filling in a 0 for each $*$ in ρ^j has exactly j nonzero blocks, as needed for (ii).

In fact, we will maintain that r^j has at least half stars in each block. Specifically, the total number of zeros in r^j will be at most the sum of the widths of the conjunctions C_1, \dots, C_j , which is at most the sum over all $i \in [j]$ of the maximum width of Φ_i , which equals $k \leq n/2$. At the end, $r^q \in (\{0,*\}^n)^{q-q}$ will be the empty tuple, which means ρ^q will be the fixed valid input $l^q = z^q$, which has all blocks nonzero.

We start with $\rho^0 = r^0 = (*^n)^q$, which indeed has no zeros. Now supposing we already have l^j and r^j satisfying all the properties from the previous two paragraphs, we explain how to obtain $l_{j+1} \in \{0,1\}^n$ and $r^{j+1} \in (\{0,*\}^n)^{q-(j+1)}$ so these properties again hold (with $l^{j+1} := l^j l_{j+1}$).

We first observe that $z^j := l^j (0^n)^{q-j}$, which is consistent with ρ^j , must be accepted by at least one conjunction from $\Phi_{j+1}, \dots, \Phi_q$. This is because z^j is already accepted by the conjunctions C_1, \dots, C_j , and these cannot be the only values of i such that $\Phi_i(z^j) = 1$ since otherwise we would have $T(z^j) = \oplus(j)$ while $\oplus_q \text{GAPOR}(z^j) = 1 - \oplus(j)$, contradicting the supposed correctness of T . Now pick any one conjunction from $\Phi_{j+1}, \dots, \Phi_q$ that accepts z^j , and assume this conjunction is C_{j+1} from Φ_{j+1} for notational convenience. Defining the partial assignment \tilde{r}^j from r^j by filling in a 0 for each $*$ corresponding to a negative literal in C_{j+1} (note that C_{j+1} has no positive literals among the last $q-j$ blocks since it accepts z^j), we see that the total number of zeros in \tilde{r}^j is at most the sum of the widths of the conjunctions C_1, \dots, C_{j+1} .

Let r^{j+1} be all but the first block of \tilde{r}^j , and obtain l_{j+1} by replacing the remaining stars in the first block of \tilde{r}^j with ones, noting that l_{j+1} is at least half ones since the first block of \tilde{r}^j was at least half stars. Now $\rho^{j+1} := l^{j+1} r^{j+1}$ maintains the desired properties: (ii) is maintained since each block of l^{j+1} is at least half ones, and r^{j+1} has at most (sum of widths of C_1, \dots, C_{j+1}) many zeros and the rest stars. To see that (i) is maintained, consider any input z consistent with ρ^{j+1} . Then z is also consistent with ρ^j and is thus accepted

by each of C_1, \dots, C_j . Moreover, $C_{j+1}(z) = 1$ since $C_{j+1}(z^j) = 1$ and z and z^j differ only in locations that are stars in \tilde{r}^j and therefore do not appear in C_{j+1} .

3.2 Proof of Theorem 1 for total functions

Fix any constant q . Let $\oplus_q \text{NONEQ}: (\{0, 1\}^n)^q \times (\{0, 1\}^n)^q \rightarrow \{0, 1\}$ be the two-party total function where Alice's and Bob's inputs are divided into q blocks $x = (x_1, \dots, x_q)$ and $y = (y_1, \dots, y_q)$ with each $x_i, y_i \in \{0, 1\}^n$, and which is defined by $\oplus_q \text{NONEQ}(x, y) := 1$ iff there are an odd number of blocks i such that $x_i \neq y_i$. Note that $\text{RP}(q)^{\text{cc}}(\oplus_q \text{NONEQ}) = O(1)$ by Lemma 8 and the standard fact that $\text{RP}^{\text{cc}}(\text{NONEQ}) = O(1)$. Thus $\oplus_q \text{NONEQ} \in \text{coRP}(q)^{\text{cc}}$. We will now prove that $\text{NP}(q)^{\text{cc}}(\oplus_q \text{NONEQ}) = \Omega(n)$.

Suppose for contradiction $\oplus_q \text{NONEQ}$ has an $\text{NP}(q)^{\text{cc}}$ protocol of cost $k \leq n/2^q$, say $\Pi = (\mathcal{R}_1, \dots, \mathcal{R}_q)$ where each \mathcal{R}_i is a nonempty set of rectangles. By Lemma 8 we may assume the protocol outputs 1 iff the input is contained in an odd number of the rectangle unions $\bigcup_{R \in \mathcal{R}_i} R$ for $i \in [q]$, in other words, $\Pi(x, y) := \oplus_q (\mathcal{R}_1(x, y), \dots, \mathcal{R}_q(x, y))$. Note that, assuming \mathcal{R}_i has cost k_i , the total number of rectangles in these unions is at most $\sum_i |\mathcal{R}_i| \leq \prod_i (|\mathcal{R}_i| + 1) \leq \prod_i 2^{k_i} = 2^k$.

We will iteratively construct a sequence of rectangles Q^j for $j = 0, \dots, q$ such that (i) there are at least j many values of i for which $Q^j \subseteq \bigcup_{R \in \mathcal{R}_i} R$, and (ii) Q^j contains an input where exactly j blocks are unequal. We obtain the contradiction when $j = q$: by (ii) some input $(x, y) \in Q^q$ has $x_i \neq y_i$ for all i and thus $\oplus_q \text{NONEQ}(x, y) = 1 - \oplus(q)$, yet by (i) we have $\mathcal{R}_i(x, y) = 1$ for all i and thus $\Pi(x, y) = \oplus(q)$, contradicting the supposed correctness of Π .

We will actually maintain stronger invariants than the above (i) and (ii): For (i), we will actually have for some j values of i —we assume they are $1, \dots, j$ for notational convenience—some individual rectangle $R_i \in \mathcal{R}_i$ contains Q^j . For (ii), Q^j will actually have the following form: for some fixed strings $a^j = (a_1, \dots, a_j) \in (\{0, 1\}^n)^j$ and $b^j = (b_1, \dots, b_j) \in (\{0, 1\}^n)^j$ such that $a_i \neq b_i$ for all $i \in [j]$, and for some nonempty set $S^j \subseteq (\{0, 1\}^n)^{q-j}$, we have $Q^j := \{a^j s : s \in S^j\} \times \{b^j s : s \in S^j\}$, which we abbreviate as $a^j S^j \times b^j S^j$. Defining a *diagonal* input in Q^j to be one of the form $(a^j s, b^j s)$ for any particular $s \in S^j$, we see that each diagonal input has exactly j unequal blocks, as needed for (ii).

In fact, we will maintain not just that S^j is nonempty, but that it is sufficiently large. Specifically, the *deficiency* of S^j , defined as $\mathbf{D}_\infty(S^j) := n(q-j) - \log |S^j|$, will be at most $(2^j - 1)(2k + 1)$. At the end, since $(2^q - 1)(2k + 1) < \infty$, this guarantees that S^q will contain at least one element from $(\{0, 1\}^n)^{q-q}$. The latter set only has one element, namely the empty tuple, so this means Q^q will contain the single input (a^q, b^q) , which has all blocks unequal.

We start with $S^0 = (\{0, 1\}^n)^q$, which indeed has $\mathbf{D}_\infty(S^0) = 0 = (2^0 - 1)(2k + 1)$, and thus Q^0 is the rectangle of all possible inputs. Now supposing we already have a^j, b^j , and S^j satisfying all the properties from the previous two paragraphs, we explain how to obtain $a_{j+1}, b_{j+1} \in \{0, 1\}^n$ and $S^{j+1} \subseteq (\{0, 1\}^n)^{q-(j+1)}$ so these properties again hold (with $a^{j+1} := a^j a_{j+1}$ and $b^{j+1} := b^j b_{j+1}$).

We first observe that each diagonal input in Q^j must be contained in at least one rectangle from $\mathcal{R}_{j+1} \cup \dots \cup \mathcal{R}_q$. This is because such an input (x, y) is already contained in the rectangles $R_1 \in \mathcal{R}_1, \dots, R_j \in \mathcal{R}_j$, and these cannot be the only values of i such that $\mathcal{R}_i(x, y) = 1$ since otherwise we would have $\Pi(x, y) = \oplus(j)$ while $\oplus_q \text{NONEQ}(x, y) = 1 - \oplus(j)$, contradicting the supposed correctness of Π . Now pick one of the (at most 2^k) rectangles from $\mathcal{R}_{j+1} \cup \dots \cup \mathcal{R}_q$ that contains the largest fraction (at least $1/2^k$) of diagonal inputs from Q^j , and assume this rectangle is $R_{j+1} \in \mathcal{R}_{j+1}$ for notational convenience. Defining $\tilde{S}^j := \{s \in S^j : (a^j s, b^j s) \in R_{j+1}\}$, we see that $\mathbf{D}_\infty(\tilde{S}^j) \leq \mathbf{D}_\infty(S^j) + k \leq (2^j - 1)(2k + 1) + k$.

Since R_{j+1} is a rectangle, it must in fact contain the entire rectangle $a^j \tilde{S}^j \times b^j \tilde{S}^j$. Since $a^j \tilde{S}^j \times b^j \tilde{S}^j \subseteq a^j S^j \times b^j S^j = Q^j$, by assumption it is also contained in each of R_1, \dots, R_j . In the end, we will ensure Q^{j+1} is a subrectangle of $a^j \tilde{S}^j \times b^j \tilde{S}^j$, which will maintain property (i): Q^{j+1} is contained in each of R_1, \dots, R_{j+1} .

To maintain (ii), we will find some $a_{j+1} \neq b_{j+1}$ and then define $S^{j+1} := \{s : a_{j+1}s \in \tilde{S}^j \text{ and } b_{j+1}s \in \tilde{S}^j\}$. Then $a_{j+1}S^{j+1} \subseteq \tilde{S}^j$ and $b_{j+1}S^{j+1} \subseteq \tilde{S}^j$ ensure that $Q^{j+1} := a^{j+1}S^{j+1} \times b^{j+1}S^{j+1}$ is indeed a subrectangle of $a^j \tilde{S}^j \times b^j \tilde{S}^j$, as we needed for (i). The fact that this can be done with a not-too-small S^{j+1} is encapsulated in the following technical lemma, which we prove shortly:

Lemma 17. *Consider any bipartite graph with left nodes U and right nodes V , and suppose $1 \geq \varepsilon \geq 2/|U|$. If an ε fraction of all possible edges are present in the graph, then there exist distinct nodes $u, u' \in U$ that have at least $(\varepsilon^2/2) \cdot |V|$ common neighbors.*

Specifically, take $U := \{0, 1\}^n$ and $V := (\{0, 1\}^n)^{q-(j+1)}$ (so $|V| = 1$ if $j = q - 1$, but that is fine), put an edge between $u \in U$ and $v \in V$ iff $uv \in \tilde{S}^j$, and let $\varepsilon := |\tilde{S}^j|/2^{n(q-j)} = 1/2^{\mathbf{D}_\infty(\tilde{S}^j)}$. Notice that $\varepsilon \geq 2/|U|$ holds since $\mathbf{D}_\infty(\tilde{S}^j) \leq (2^j - 1)(2k + 1) + k \leq 2^{j+1}k - 1 \leq n - 1$ follows from our assumption that $k \leq n/2^q$. Thus [Lemma 17](#) guarantees we can pick strings $a_{j+1} \neq b_{j+1}$ (corresponding to the nodes u, u') such that S^{j+1} (the set of common neighbors) has size at least $(\varepsilon^2/2) \cdot 2^{n(q-(j+1))}$. Thus

$$\begin{aligned} \mathbf{D}_\infty(S^{j+1}) &:= n(q - (j + 1)) - \log |S^{j+1}| \leq \log(2/\varepsilon^2) = 2\mathbf{D}_\infty(\tilde{S}^j) + 1 \\ &\leq 2((2^j - 1)(2k + 1) + k) + 1 = (2(2^j - 1) + 1)(2k + 1) = (2^{j+1} - 1)(2k + 1) \end{aligned}$$

as we needed for (ii). This finishes the proof of [Theorem 1](#).

Proof of Lemma 17. Let d_u and d_v denote the degrees of nodes $u \in U$ and $v \in V$, and let $d_{u,u'}$ denote the number of common neighbors of $u, u' \in U$. Summing over ordered pairs u, u' of not-necessarily-distinct left nodes, we have

$$\sum_{u, u' \in U} d_{u, u'} = \sum_{v \in V} d_v^2 \geq \left(\sum_{v \in V} d_v \right)^2 / |V| = \varepsilon^2 \cdot |U|^2 \cdot |V|$$

by Cauchy–Schwarz and the assumption $\sum_{v \in V} d_v = \varepsilon \cdot |U| \cdot |V|$. Now sampling u, u' independently uniformly at random from U , we have

$$\varepsilon^2 \cdot |V| \leq \mathbb{E}_{u, u'}[d_{u, u'}] \leq \mathbb{E}_{u, u'}[d_{u, u'} \mid u \neq u'] + \mathbb{E}_u[d_u] \cdot \Pr_{u, u'}[u = u']$$

(the conditioning is valid by the assumption $|U| \geq 2$). Since $\mathbb{E}_u[d_u] = \varepsilon \cdot |V|$ and $\Pr_{u, u'}[u = u'] = 1/|U|$, rearranging gives

$$\mathbb{E}_{u, u'}[d_{u, u'} \mid u \neq u'] \geq \varepsilon^2 \cdot |V| - \varepsilon \cdot |V|/|U| \geq (\varepsilon^2/2) \cdot |V|$$

where the last inequality holds by the assumption $1/|U| \leq \varepsilon/2$. Thus there must be some $u \neq u'$ such that $d_{u, u'}$ is at least this large. \square

3.3 Proof of [Theorem 3](#)

Restatement of [Theorem 3](#). For partial functions, $\text{RP}(q + 1)^{\text{cc}} \cap \text{coRP}(q + 1)^{\text{cc}} \not\subseteq \mathbb{P}_{\parallel}^{\text{NP}[q]^{\text{cc}}}$ for every constant q .

Fix any constant q . Let $\text{WHICH}_{\oplus_{q+1}\text{GAPOR}}: (\{0, 1\}^{2n})^{2(q+1)} \rightarrow \{0, 1\}$ be the following partial function: The input is divided into two halves $z = (z^0, z^1)$, and each half is divided into $q + 1$ blocks $z^h = (z_1^h, \dots, z_{q+1}^h)$ having the promise that each $z_i^h \in \{0, 1\}^{2n}$ is either all zeros or at least a *quarter* ones, and moreover it is promised that the number of nonzero blocks in z^0 has the opposite parity as the number of nonzero blocks in z^1 (call such an input *valid*). The partial function is defined by

$$\text{WHICH}_{\oplus_{q+1}\text{GAPOR}}(z) = \begin{cases} 1 & \text{if the number of nonzero blocks is odd in } z^0 \text{ and even in } z^1 \\ 0 & \text{if the number of nonzero blocks is even in } z^0 \text{ and odd in } z^1 \end{cases}$$

We henceforth abbreviate $\text{WHICH}_{\oplus_{q+1}\text{GAPOR}}$ as f . Note that $\text{RP}(q + 1)^{\text{dt}}(f) = O(1)$ by applying the $\text{RP}(q + 1)^{\text{dt}}$ decision tree for $\oplus_{q+1}\text{GAPOR}$ on z^0 (adapted for the different block length and different threshold for fraction of ones in a block). By symmetry (focusing on z^1), we also have $\text{RP}(q + 1)^{\text{dt}}(\bar{f}) = O(1)$. Letting $g: [m] \times \{0, 1\}^m \rightarrow \{0, 1\}$ be the index gadget with $m := N^{20}$ where $N := 4(q + 1)n$, this implies that

$$\text{RP}(q + 1)^{\text{cc}}(f \circ g^N) = O(\log n) \quad \text{and} \quad \text{coRP}(q + 1)^{\text{cc}}(f \circ g^N) = O(\log n)$$

(by the “easy direction” of $\text{RP}(q + 1)$ lifting) and thus $f \circ g^N \in \text{RP}(q + 1)^{\text{cc}} \cap \text{coRP}(q + 1)^{\text{cc}}$. We will now prove that $\mathbb{P}_{\parallel}^{\text{NP}[q]^{\text{dt}}}(f) = \Omega(n)$, which by [Theorem 4.\(ii\)](#) implies that $\mathbb{P}_{\parallel}^{\text{NP}[q]^{\text{cc}}}(f \circ g^N) = \Omega(n \log n)$.

We show this by reduction from the fact that $\text{NP}(q)^{\text{dt}}(\oplus_q\text{GAPOR}) = \Omega(n)$. We henceforth abbreviate $\oplus_q\text{GAPOR}$ as f' , and as in [Section 3.1](#) we assume f' has block length n and threshold half. Supposing f has a $\mathbb{P}_{\parallel}^{\text{NP}[q]^{\text{dt}}}$ decision tree T of cost $k \leq n/2$, say with deterministic phase T_{det} , we will use it to construct an $\text{NP}(q)^{\text{dt}}$ decision tree T' of cost at most k for f' .

By [Lemma 16](#) we may assume that each leaf of T_{det} produces a single $\text{NP}(q)^{\text{dt}}$ decision tree and chooses whether to output the same or opposite answer as that decision tree. Follow the root-to-leaf path in T_{det} where all queries are answered with zero. Let $\rho \in (\{0, *\}^{2n})^{2(q+1)}$ be the partial assignment with at most $k \leq n/2$ zeros that records these queries (so an input leads to this leaf iff it is consistent with ρ). Let $T_{\text{leaf}} = (\Phi_1, \dots, \Phi_q)$ be the $\text{NP}(q)^{\text{dt}}$ decision tree of cost at most k produced at this leaf, where each Φ_i is a DNF. By symmetry, we assume (without loss of generality) that this leaf chooses to output the same answer as T_{leaf} .

Given any valid input z' to f' , we show how to map it to a valid input z to f such that (i) $f'(z') = f(z)$, (ii) z is consistent with ρ , and (iii) each bit of z either is fixed or is some preselected bit of z' . Since (iii) implies that $T_{\text{leaf}}(z)$ can be viewed as an $\text{NP}(q)^{\text{dt}}$ decision tree $T'(z')$ by substituting a constant or variable of z' in for each variable of z (which does not increase the width of any conjunction), and T' would correctly compute f' since

$$f'(z') = f(z) = T(z) = T_{\text{leaf}}(z) = T'(z')$$

by (i), correctness of T , (ii), and (iii) respectively, this would show that $\text{NP}(q)^{\text{dt}}(f') \leq k \leq n/2$, which we know is false from [Section 3.1](#).

To define z , we start with ρ (that is, we place zeros everywhere ρ requires, thus ensuring (ii)). Since ρ has at most n zeros in each block (indeed, at most $n/2$ zeros total), we can then place more zeros in such a way that each block now has exactly n zeros and n stars. Next we replace the stars in z_{q+1}^0 with ones and replace the stars in z_{q+1}^1 with zeros. Finally, for each $i \in [q]$, we fill in the stars of z_i^0 with a copy of z'_i and fill in the stars of z_i^1 with another copy of z'_i . This construction satisfies (iii).

To verify (i), first observe that since each block of z' is either all zeros or at least half ($n/2$) ones, this ensures each block of z is either all zeros or at least a quarter ($2n/4$) ones. Furthermore, if z' has exactly ℓ nonzero blocks then the number of nonzero blocks is $\ell + 1$ in z^0 (since z_{q+1}^0 is nonzero) and ℓ in z^1 (since z_{q+1}^1 is all zeros). Hence if $f'(z') = 1$ (ℓ is even) then $f(z) = 1$ (since $\ell + 1$ is odd and ℓ is even), and if $f'(z') = 0$ (ℓ is odd) then $f(z) = 0$ (since $\ell + 1$ is even and ℓ is odd). Thus $f'(z') = f(z)$, and this finishes the proof of [Theorem 3](#).

4 Total function collapse

Restatement of [Theorem 2](#). *For total functions,*

$$\mathbb{P}_{\parallel}^{\text{NP}[q]^{\text{cc}}} = \text{NP}(q+1)^{\text{cc}} \cap \text{coNP}(q+1)^{\text{cc}} \quad \text{and} \quad \mathbb{P}_{\parallel}^{\text{RP}[q]^{\text{cc}}} = \text{RP}(q+1)^{\text{cc}} \cap \text{coRP}(q+1)^{\text{cc}}$$

for every constant q .

We start with the intuition for the nondeterministic case of [Theorem 2](#). This is proved in a way similar to the specific result for $q = 0$ (that is, for total functions, $\mathbb{P}^{\text{cc}} = \text{NP}^{\text{cc}} \cap \text{coNP}^{\text{cc}}$). In that proof, Alice and Bob can use the fact that the rectangles in the NP^{cc} protocol's 1-monochromatic covering of F are disjoint from the rectangles in the coNP^{cc} protocol's 0-monochromatic covering. Specifically, if $F(x, y) = 1$, then (x, y) is in some 1-rectangle, which is row-disjoint or column-disjoint from each 0-rectangle. (If a 1-rectangle and 0-rectangle shared a row and a column, they would intersect, which is not possible for a total function.) Therefore, Alice and Bob can repeatedly eliminate from consideration at least half of the remaining 0-rectangles, by identifying a 1-rectangle that either has x in its row set but is row-disjoint from at least half the remaining 0-rectangles, or has y in its column set but is column-disjoint from at least half the remaining 0-rectangles. If (x, y) is indeed in a 1-rectangle, then this process can always continue until there are no 0-rectangles left. If (x, y) is in a 0-rectangle, then this process will never eliminate that rectangle, so the process will halt with a nonempty set of 0-rectangles.

We repeat a similar argument, but using the “top level” of the $\text{NP}(q+1)^{\text{cc}}$ and $\text{coNP}(q+1)^{\text{cc}}$ protocols for F as our monochromatic rectangle sets. Here we think of a $\text{coNP}(q+1)^{\text{cc}}$ protocol as computing F by applying $\overline{\Delta}_{q+1}$ (with negations pushed to the leaves) to the indicators for $q+1$ rectangle unions. Depending on the parity of q , the rectangle union \mathcal{R}_{q+1} queried at depth 1 of the $\text{NP}(q+1)^{\text{cc}}$ protocol will correspond to either 1-monochromatic rectangles or 0-monochromatic rectangles for F . The rectangle union \mathcal{R}'_{q+1} queried at depth 1 of the $\text{coNP}(q+1)^{\text{cc}}$ protocol will be the opposite color of monochromatic rectangles. Crucially,

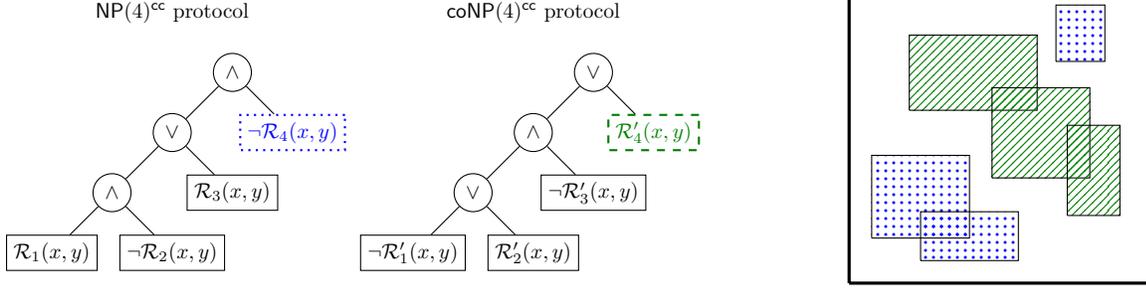


Figure 3: If a total function has an $\text{NP}(4)^{\text{cc}}$ protocol and a $\text{coNP}(4)^{\text{cc}}$ protocol, then the rectangle unions from the NP^{cc} functions at depth one of each protocol are disjoint.

this means that no input is in a rectangle from both of these sets (as we are assuming F is total). See Figure 3 for an illustration.

A key observation is that a deterministic protocol similar to the one used in the $q = 0$ case, ran using these top-level rectangle sets, will return the correct answer *under the promise that (x, y) is in one of these rectangles*. Say, for example, that (x, y) is in some rectangle in the 1-monochromatic top-level set. Then the deterministic protocol will successfully eliminate all 0-rectangles from the other top-level set, and will announce that the answer is 1. If (x, y) was in one of the 0-rectangles, then that rectangle will never be eliminated, and so the protocol would announce that the answer is 0.

If (x, y) is in the top-level rectangle union for one of the protocols, then (x, y) is not in the top-level rectangle union of the *other* protocol, so $F(x, y)$ can be computed by the other protocol but where the top level is skipped (resulting in only q many NP^{cc} oracle queries). This boils down to the observation that $\Delta_{q+1}(z_1, \dots, z_q, 0) = \Delta_q(z_1, \dots, z_q)$.

What if (x, y) is in neither top-level rectangle union? Then we can make no guarantees about the behavior of the deterministic protocol—it might answer 0 or 1 (which we interpret as merely a “guess” for $F(x, y)$). However, in this case *both* protocols correctly compute $F(x, y)$ even if the top level is skipped. Therefore, we will still get the correct answer no matter which guess is produced by the deterministic protocol.

What follows is the formal proof.

Proof of Theorem 2 for the nondeterministic case. We already know $\text{P}_{\parallel}^{\text{NP}[q]^{\text{cc}}} \subseteq \text{NP}(q+1)^{\text{cc}} \cap \text{coNP}(q+1)^{\text{cc}}$ [HR90]. Let $F: \mathcal{X} \times \mathcal{Y} \rightarrow \{0, 1\}$ be a two-party total function with $\max\{\text{NP}(q+1)^{\text{cc}}(F), \text{coNP}(q+1)^{\text{cc}}(F)\} = k$. Consider any $\text{NP}(q+1)^{\text{cc}}$ protocol for F with cost at most k , say $\Pi = (\mathcal{R}_1, \dots, \mathcal{R}_{q+1})$ where each \mathcal{R}_i is a nonempty set of rectangles. Consider any $\text{coNP}(q+1)^{\text{cc}}$ protocol for F with cost at most k , say $\Pi' = (\mathcal{R}'_1, \dots, \mathcal{R}'_{q+1})$ meaning that $F(x, y) = \overline{\Delta}_{q+1}(\mathcal{R}'_1(x, y), \dots, \mathcal{R}'_{q+1}(x, y))$.

Out of the two protocols Π and Π' , let $\Pi^\wedge = (\mathcal{R}_1^\wedge, \dots, \mathcal{R}_{q+1}^\wedge)$ be whichever one has \wedge as its root gate, and let $\Pi^\vee = (\mathcal{R}_1^\vee, \dots, \mathcal{R}_{q+1}^\vee)$ be whichever one has \vee as its root gate (after pushing negations to the leaves in Π'). Similarly, out of the two functions Δ_q and $\overline{\Delta}_q$, let Δ_q^\wedge be whichever one has \wedge as its root gate, and let Δ_q^\vee be whichever one has \vee as its root gate. In other words:

	Π^\wedge	Π^\vee	Δ_q^\wedge	Δ_q^\vee
if q is odd	Π	Π'	$\overline{\Delta}_q$	Δ_q
if q is even	Π'	Π	Δ_q	$\overline{\Delta}_q$

We claim that Algorithm 1 is a $\text{P}_{\parallel}^{\text{NP}[q]^{\text{cc}}}$ protocol with cost $O(k^2)$ that correctly computes F .

Claim 18. *Algorithm 1 is a $\text{P}_{\parallel}^{\text{NP}[q]^{\text{cc}}}$ protocol with cost $O(k^2)$.*

Proof of Claim 18. GUESS is a deterministic protocol. In each iteration of the loop, either at least half of the remaining rectangles in \mathcal{R}^0 get removed, or the loop terminates. Since $|\mathcal{R}^0| \leq 2^k$ at the beginning, there can be at most $k+1$ iterations. Each iteration involves up to $k+1$ bits of communication, to specify one of the at most 2^k rectangles in \mathcal{R}^1 (or indicate the non-existence of a suitable rectangle in \mathcal{R}^1). Thus

Algorithm 1 $\mathsf{P}_{\parallel}^{\text{NP}[q]\text{cc}}$ protocol for F

```

1: function GUESS
2:    $\mathcal{R}^0 \leftarrow \mathcal{R}_{q+1}^{\wedge}$  (set of 0-monochromatic rectangles)
3:    $\mathcal{R}^1 \leftarrow \mathcal{R}_{q+1}^{\vee}$  (set of 1-monochromatic rectangles)
4:   loop
5:     if  $\mathcal{R}_0 = \emptyset$  then return 1
6:     else if there exists an  $R \in \mathcal{R}_1$  whose row set contains  $x$ 
       and which is row-disjoint from at least half of the rectangles in  $\mathcal{R}_0$  then
7:       Alice announces such an  $R$ 
8:       remove from  $\mathcal{R}_0$  all rectangles that are row-disjoint from  $R$ 
9:     else if there exists an  $R \in \mathcal{R}_1$  whose column set contains  $y$ 
       and which is column-disjoint from at least half of the rectangles in  $\mathcal{R}_0$  then
10:      Bob announces such an  $R$ 
11:      remove from  $\mathcal{R}_0$  all rectangles that are column-disjoint from  $R$ 
12:     else return 0

13: function MAIN
14:   if GUESS = 0 then return  $\Delta_q^{\wedge}(\mathcal{R}_1^{\vee}(x, y), \dots, \mathcal{R}_q^{\vee}(x, y))$ 
15:   if GUESS = 1 then return  $\Delta_q^{\vee}(\mathcal{R}_1^{\wedge}(x, y), \dots, \mathcal{R}_q^{\wedge}(x, y))$ 

```

GUESS has communication cost $O(k^2)$. Also, the q many nonadaptive NP^{cc} oracle queries on line 14 or line 15 contribute at most k to the cost since they are just part of Π or Π' , so the overall cost is still $O(k^2)$. \square

Claim 19. *If $\mathcal{R}_{q+1}^{\wedge}(x, y) = 1$ then GUESS = 0. If $\mathcal{R}_{q+1}^{\vee}(x, y) = 1$ then GUESS = 1.*

Proof of Claim 19. Assume $\mathcal{R}_{q+1}^{\wedge}(x, y) = 1$, so $(x, y) \in R$ for some $R \in \mathcal{R}_{q+1}^{\wedge}$. Note that GUESS only removes a rectangle from \mathcal{R}^0 when it is certain that (x, y) is not in that rectangle (because x is not in its row set or y is not in its column set). Thus R will never be removed from \mathcal{R}^0 throughout GUESS, which means \mathcal{R}^0 will never be empty, and GUESS will eventually return 0.

Assume $\mathcal{R}_{q+1}^{\vee}(x, y) = 1$, so $(x, y) \in R'$ for some $R' \in \mathcal{R}_{q+1}^{\vee}$. Observe that each rectangle in $\mathcal{R}_{q+1}^{\wedge}$ is 0-monochromatic for F (since Π^{\wedge} outputs 0 on such inputs) and each rectangle in \mathcal{R}_{q+1}^{\vee} is 1-monochromatic for F (since Π^{\vee} outputs 1 on such inputs). Since F is total, R' is disjoint—and hence either row-disjoint or column-disjoint—from each rectangle $R \in \mathcal{R}_{q+1}^{\wedge}$ (since R' contains only 1-inputs and R contains only 0-inputs). Thus in each iteration of the loop, either at least half the remaining rectangles in \mathcal{R}^0 are row-disjoint from $R' \in \mathcal{R}^1$ (which Alice would notice) or at least half are column-disjoint from R' (which Bob would notice). Either way, the loop will continue with one party announcing a rectangle (not necessarily R') and shrinking \mathcal{R}^0 , which means the loop will not halt until $\mathcal{R}^0 = \emptyset$, and GUESS will return 1. \square

Claim 20. *For all $x \in \mathcal{X}$ and $y \in \mathcal{Y}$, Algorithm 1 correctly computes $F(x, y)$.*

Proof of Claim 20. If GUESS = 0 then $\mathcal{R}_{q+1}^{\vee}(x, y) = 0$ by Claim 19, and thus $F(x, y) = \Pi^{\vee}(x, y) = \Delta_q^{\wedge}(\mathcal{R}_1^{\vee}(x, y), \dots, \mathcal{R}_q^{\vee}(x, y)) \vee \mathcal{R}_{q+1}^{\vee}(x, y)$ equals the output on line 14. If GUESS = 1 then $\mathcal{R}_{q+1}^{\wedge}(x, y) = 0$ by Claim 19, and thus $F(x, y) = \Pi^{\wedge}(x, y) = \Delta_q^{\vee}(\mathcal{R}_1^{\wedge}(x, y), \dots, \mathcal{R}_q^{\wedge}(x, y)) \wedge \neg \mathcal{R}_{q+1}^{\wedge}(x, y)$ equals the output on line 15. \square

Together, Claim 18 and Claim 20 show that Algorithm 1 witnesses $\mathsf{P}_{\parallel}^{\text{NP}[q]\text{cc}}(F) \leq O(k^2)$. This completes the proof that for total functions, $\mathsf{P}_{\parallel}^{\text{NP}[q]\text{cc}} = \text{NP}(q+1)^{\text{cc}} \cap \text{coNP}(q+1)^{\text{cc}}$. \square

The corresponding argument for the Randomized Boolean Hierarchy is very similar. Note that since $\text{RP}^{\text{cc}} \subseteq \text{NP}^{\text{cc}}$, we can simply interpret the $(q+1)^{\text{st}}$ components of our $\text{RP}(q+1)^{\text{cc}}$ and $\text{coRP}(q+1)^{\text{cc}}$ protocols as NP^{cc} protocols. Then Algorithm 1 would work in exactly the same way, except instead of using $\mathcal{R}_1^{\vee}, \dots, \mathcal{R}_q^{\vee}$ or $\mathcal{R}_1^{\wedge}, \dots, \mathcal{R}_q^{\wedge}$ as the oracle queries at the end, these would be replaced by RP^{cc} protocols (the first q components of our $\text{RP}(q+1)^{\text{cc}}$ or $\text{coRP}(q+1)^{\text{cc}}$ protocols).

We remark that a similar approach shows that $P_{\parallel}^{\text{NP}[q]^{\text{dt}}} = \text{NP}(q+1)^{\text{dt}} \cap \text{coNP}(q+1)^{\text{dt}}$ for total functions. (The corresponding result for randomized query complexity follows anyway from $P^{\text{dt}} = \text{BPP}^{\text{dt}}$ [Nis91].)

5 Query-to-communication lifting for $\text{NP}(q)$

Restatement of Theorem 4.(i). *For every partial function $f: \{0,1\}^n \rightarrow \{0,1\}$ and every constant q ,*

$$\text{NP}(q)^{\text{cc}}(f \circ g^n) = \text{NP}(q)^{\text{dt}}(f) \cdot \Theta(\log n)$$

where $g: [m] \times \{0,1\}^m \rightarrow \{0,1\}$ is the index gadget defined by $g(x,y) = y_x$ with $m := n^{20}$.

The big- O direction follows immediately from the same fact for NP : for every f , $\text{NP}^{\text{cc}}(f \circ g^n) = \text{NP}^{\text{dt}}(f) \cdot O(\log n)$ holds by replacing each of the $n^{O(k)}$ conjunctions in a width- k DNF with m^k rectangles (each of which contains inputs where the gadget outputs satisfy the conjunction), for a total of $n^{O(k)}m^k = 2^{k \cdot O(\log n)}$ rectangles. In the rest of this section we prove the big- Ω direction. By Lemma 13 it suffices to show

$$\text{DL}(q)^{\text{cc}}(f \circ g^n) = \text{DL}(q)^{\text{dt}}(f) \cdot \Omega(\log n).$$

5.1 Technical preliminaries

Our proof is closely related to the P^{NP} lifting theorem of Göös, Kamath, Pitassi, and Watson [GKPW17], so we start by recalling some definitions and lemmas that were used in that work. We will need to tweak some of the statements and parameters, though.

Define $G: [m]^n \times (\{0,1\}^m)^n \rightarrow \{0,1\}^n$ as $G := g^n$. This partitions the input domain into 2^n slices $G^{-1}(z) = \{(x,y) : g(x_i, y_i) = z_i \text{ for all } i \in [n]\}$, one for each $z \in \{0,1\}^n$. For a set $Z \subseteq \{0,1\}^n$, let $G^{-1}(Z) := \bigcup_{z \in Z} G^{-1}(z)$.

Consider sets $A \subseteq [m]^n$ and $B \subseteq (\{0,1\}^m)^n$. For $I \subseteq [n]$, we let $A_I := \{x_I : x \in A\}$ and $B_I := \{y_I : y \in B\}$ be the *projections* onto the coordinates of I . The *min-entropy* of a random variable \mathbf{x} is $\mathbf{H}_{\infty}(\mathbf{x}) := \min_x \log(1/\Pr[\mathbf{x} = x])$. We say A is δ -dense if the uniform random variable \mathbf{x} over A satisfies the following: for every nonempty $I \subseteq [n]$, $\mathbf{H}_{\infty}(\mathbf{x}_I) \geq \delta|I| \log m$ (that is, the min-entropy of the marginal distribution of \mathbf{x} on coordinates I is at least a δ fraction of the maximum possible for a distribution over $[m]^I$). The *deficiency* of B is $\mathbf{D}_{\infty}(B) := mn - \log |B|$.

Lemma 21 ([GKPW17, Lemma 11]). *If $A \subseteq [m]^n$ is 0.8-dense and $B \subseteq (\{0,1\}^m)^n$ has deficiency at most n^4 , then $G(A \times B) = \{0,1\}^n$, that is, for every $z \in \{0,1\}^n$ there are $x \in A$ and $y \in B$ with $G(x,y) = z$.*

Here the density parameter is $\delta = 0.8$ and the deficiency is $\mathbf{D}_{\infty}(B) \leq n^4$, instead of $\delta = 0.9$ and $\mathbf{D}_{\infty}(B) \leq n^2$ as in the original. Lemma 21 still holds because our gadget size has increased: we use $m := n^{20}$, whereas [GKPW17] used $m := n^4$. This can be verified by a simple substitution in the proof.

The next lemma is altered enough from the original that we will reprove it here.

Lemma 22 (A more general version of [GKPW17, Claim 12]). *Let $\mathcal{X} \subseteq [m]^n$ be 0.85-dense. If $A' \subseteq \mathcal{X}$ satisfies $|A'| \geq |\mathcal{X}|/2^{k+1}$ then there exist an $I \subseteq [n]$ of size $|I| < 20(k+1)/\log m$ and an $A \subseteq A'$ such that A is fixed on coordinates I and 0.8-dense on all other coordinates.*

The original version was for the special case $\mathcal{X} = [m]^n$. We observe that it is sufficient for \mathcal{X} to be δ -dense, for some suitable constant δ greater than the desired density of the resulting set A (here we use $\delta = 0.85$).

Proof of Lemma 22. If A' is already 0.8-dense, then we can simply take $I := \emptyset$ and $A := A'$, so assume otherwise. Let $I \subseteq [n]$ be a maximal set of coordinates that violates 0.8-density: for the uniform random variable \mathbf{x} over A' , $\mathbf{H}_{\infty}(\mathbf{x}_I) < 0.8|I| \log m$. Since the uniform random variable \mathbf{X} over \mathcal{X} has $\mathbf{H}_{\infty}(\mathbf{X}_I) \geq 0.85|I| \log m$, and $|A'| \geq |\mathcal{X}|/2^{k+1}$, we also have $\mathbf{H}_{\infty}(\mathbf{x}_I) \geq 0.85|I| \log m - (k+1)$. Combining these, we get that $k+1 > 0.05|I| \log m$, which implies that $|I| < 20(k+1)/\log m$.

We can now simply choose some $\alpha \in [m]^I$ such that $\Pr[\mathbf{x}_I = \alpha] > 2^{0.8|I| \log m}$, and take $A := \{x \in A' : x_I = \alpha\}$. This certainly satisfies that A is fixed on coordinates I . To see that A is 0.8-dense on all other coordinates, assume for contradiction there is some nonempty $J \subseteq [n] \setminus I$ witnessing a 0.8-density violation of A . Then it is straightforward to check that $I \cup J$ would be a set of coordinates witnessing a 0.8-density violation of A' , which contradicts the maximality of I . \square

5.2 The simulation

We exhibit an algorithm that takes a q -alternating rectangle decision list \mathcal{L}_R for $f \circ g^n$ of cost k , and converts it to a q -alternating conjunction decision list \mathcal{L}_C for f of cost $O(k/\log n)$. The argument from [GKPW17] does exactly this except without preserving the bound on the number of alternations. In [GKPW17] the argument is formulated using a “dual” characterization of DL^{dt} , but it has the effect of building \mathcal{L}_C in order, obtaining each conjunction by “extracting” it from one of the rectangles in \mathcal{L}_R . The trouble is that the rectangles are not necessarily “extracted from” in order: after extracting a conjunction from some rectangle, the *next* conjunction that gets put in \mathcal{L}_C may be extracted from a rectangle that is *earlier* in \mathcal{L}_R . Thus \mathcal{L}_C may end up with more alternations than \mathcal{L}_R .

To fix this, we convert the argument to a “primal” form and argue that it still works when we force the rectangles to be extracted from in order. The high-level view is that we iterate through the rectangles of \mathcal{L}_R in order, and for each we extract as many conjunctions as we can until the rectangle becomes “exhausted”, at which time we remove the remaining “error” portion of the rectangle (by deleting few rows and columns) and move on to the next rectangle. With this modification, the rest of the technical details from [GKPW17] continue to work, and it now preserves the number of alternations.

At any step of this process, we let $X \times Y$ be the remaining rows and columns (after having removed the error portion of all previous rectangles in \mathcal{L}_R), and we let $Z \subseteq \{0, 1\}^n$ be the remaining inputs to f (which have not been accepted by any previous conjunctions we put in \mathcal{L}_C). Suppose (R_i, ℓ_i) is our current entry in \mathcal{L}_R . The goal is to find a subrectangle $A \times B \subseteq R_i \cap (X \times Y)$ that is “conjunction-like” in the sense that $G(A \times B)$ is exactly the inputs accepted by some small-width conjunction C , and such that among all remaining inputs $z \in Z$, C only accepts those with $f(z) = \oplus(\ell_i)$. These properties would ensure it is safe to put (C, ℓ_i) next in \mathcal{L}_C .

Combining Lemma 21 and Lemma 22 (using $\mathcal{X} = [m]^n$) suggests an approach for finding a conjunction-like subrectangle: If A' is not too small, we can restrict it to A that is fixed on few coordinates I and dense on the rest (by Lemma 22). If B is also not too small (low deficiency) and fixed on coordinates I , then $G(A \times B)$ is fixed on I and takes on all possible values on the remaining coordinates (by Lemma 21, which still works with $[n] \setminus I$ in place of $[n]$). In other words, $G(A \times B) = C^{-1}(1)$ for a small-width conjunction C , as desired.

The other property we needed to ensure is that if this C is the first conjunction in \mathcal{L}_C to accept a particular z , then $f(z) = \oplus(\ell_i)$ (so \mathcal{L}_C is correct). This will follow if we know there is some $(x, y) \in G^{-1}(z)$ such that R_i is the first rectangle in \mathcal{L}_R to contain (x, y) , as that guarantees $f(z) = f(G(x, y)) = (f \circ g^n)(x, y) = \oplus(\ell_i)$. It turns out this will hold automatically if $A \times B \subseteq R_i \cap (X \times Y)$, because $A \times B$ touches the slice of every z that is accepted by C , and all inputs $(x, y) \in G^{-1}(Z)$ that were in some R_j with $j < i$ have already been removed from $X \times Y$.

Our algorithm for building \mathcal{L}_C from \mathcal{L}_R is shown in Algorithm 2. It is described as starting from some arbitrary initial rectangle $\mathcal{X} \times \mathcal{Y}$. For the purpose of proving Theorem 4.(i) we only need to take $\mathcal{X} = [m]^n$ and $\mathcal{Y} = (\{0, 1\}^m)^n$, but when we invoke this as a component in the proof of Theorem 4.(ii) we will need to start from some $\mathcal{X} \times \mathcal{Y}$ that is merely “dense \times large” rather than the full input domain, so we state this more general version now.

Lemma 23. *If \mathcal{L}_R computes $f \circ g^n$ on $\mathcal{X} \times \mathcal{Y}$ and has cost k , and if \mathcal{X} is 0.85-dense and $\mathbf{D}_\infty(\mathcal{Y}) \leq n^3$, then \mathcal{L}_C produced by Algorithm 2 computes f and has cost $O(k/\log n)$. Moreover, if \mathcal{L}_R is q -alternating then so is \mathcal{L}_C .*

Proof. To verify the cost, just note that lines 11 and 12 always succeed by Lemma 22 (since \mathcal{X} is 0.85-dense and $|A'| \geq |\mathcal{X}|/2^{k+1}$), so when a conjunction is added to \mathcal{L}_C on lines 14 and 15, it has width $|I| < 20(k+1)/\log m = O(k/\log n)$. On line 13 we have $|B| \geq |Y_{x'}|/2^{m|I|} \geq 2^{mn-n^4-m|I|} = 2^{m(n-|I|)-n^4}$ (since $x' \in A \subseteq A'$) and therefore $\mathbf{D}_\infty(B_{[n] \setminus I}) \leq n^4$ (relative to $(\{0, 1\}^m)^{[n] \setminus I}$). Thus by applying Lemma 21 to $A_{[m] \setminus I}$ (which is 0.8-dense) and $B_{[n] \setminus I}$ we have $g^{[m] \setminus I}(A_{[m] \setminus I} \times B_{[n] \setminus I}) = \{0, 1\}^{n-|I|}$ and therefore $G(A \times B) = C^{-1}(1)$. (Lemma 21 works with the same parameters even though the sets are now on fewer than n coordinates.)

The algorithm terminates because Z always shrinks on line 16: for any $y \in B$ we have $G(x', y) \in Z$ (from the definition of $Y_{x'}$) and $C(G(x', y)) = 1$ (since $x'_I = \alpha$ and $y_I = \beta$ and thus $G(x', y)_I = g^I(\alpha, \beta)$).

Algorithm 2 Simulation algorithm

In: $\mathcal{L}_R = (R_1, \ell_1), \dots, (R_{2^k}, \ell_{2^k})$ and $\mathcal{X} \subseteq [m]^n$, $\mathcal{Y} \subseteq (\{0, 1\}^m)^n$
Out: \mathcal{L}_C

- 1: initialize $X \leftarrow \mathcal{X}$, $Y \leftarrow \mathcal{Y}$, $Z \leftarrow \text{domain of } f$, $\mathcal{L}_C \leftarrow \text{empty list}$
- 2: **for** $i = 1$ to 2^k **do**
- 3: **while** $Z \neq \emptyset$ **do**
- 4: for each $x \in X$, let $Y_x := \{y \in Y : (x, y) \in R_i \cap G^{-1}(Z)\}$
- 5: let $A' := \{x \in X : |Y_x| \geq 2^{mn-n^4}\}$
- 6: **if** $|A'| \leq |\mathcal{X}|/2^{k+1}$ **then**
- 7: update $X \leftarrow X \setminus A'$
- 8: update $Y \leftarrow Y \setminus \bigcup_{x \in X \setminus A'} Y_x$
- 9: break out of inner loop
- 10: **else** $|A'| > |\mathcal{X}|/2^{k+1}$
- 11: let $A \subseteq A'$, $I \subseteq [n]$, $\alpha \in [m]^I$ be such that:
- 12: $|I| = O(k/\log n)$, A_I is fixed to α , and $A_{[n] \setminus I}$ is 0.8-dense
- 13: pick any $x' \in A$ and choose $\beta \in (\{0, 1\}^m)^I$ to maximize the size of $B := \{y \in Y_{x'} : y_I = \beta\}$
- 14: let C be the conjunction “ $z_I = g^I(\alpha, \beta)$ ”
- 15: update \mathcal{L}_C by appending (C, ℓ_i) to it
- 16: update $Z \leftarrow Z \setminus C^{-1}(1)$

The algorithm maintains the invariant that for all $j < i$, $R_j \cap (X \times Y) \cap G^{-1}(Z) = \emptyset$. This vacuously holds at the beginning, and is clearly maintained in the **else** case because i stays the same and nothing gets added to X , Y , or Z . Lines 7 and 8 maintain the invariant in the **if** case because the removed rows and columns cover all of $R_i \cap (X \times Y) \cap G^{-1}(Z)$ and i goes up by 1.

Next we argue that when the algorithm terminates, Z must be empty. In each iteration of the outer loop, we throw out at most $|\mathcal{X}|/2^{k+1}$ rows and at most $|\mathcal{X}| \cdot 2^{mn-n^4} \leq m^n \cdot 2^{mn-n^4} \leq 2^{mn-n^3}/2^{k+1} \leq |\mathcal{Y}|/2^{k+1}$ columns. (We throw out columns in Y_x for $x \notin A'$, all of these Y_x had the property $|Y_x| < 2^{mn-n^4}$, we do this for at most $|\mathcal{X}|$ values of x , and $n^4 - n \log m \geq n^3 + k + 1$.) Since the outer loop executes 2^k times, by the end at most half the rows of \mathcal{X} and half the columns of \mathcal{Y} have been discarded, so $|X| \geq |\mathcal{X}|/2$ and $|Y| \geq |\mathcal{Y}|/2$. This means X is essentially as dense as \mathcal{X} (only a -1 loss in any $\mathbf{H}_\infty(\mathbf{x}_I)$) and Y is essentially as low-deficiency as \mathcal{Y} (only a $+1$ loss in \mathbf{D}_∞). Thus Lemma 21 (with a tiny perturbation of the parameters, which does not affect the result) shows that $G(X \times Y) = \{0, 1\}^n$. However, the last rectangle that is processed, R_{2^k} , contains all of $\mathcal{X} \times \mathcal{Y}$ by definition (since we assume \mathcal{L}_R is correct on $\mathcal{X} \times \mathcal{Y}$). So, the invariant guarantees $(X \times Y) \cap G^{-1}(Z) = \emptyset$ at termination. This can only happen if $G^{-1}(Z) = \emptyset$ and thus $Z = \emptyset$ (since $G(X \times Y) = \{0, 1\}^n$).

We now argue that \mathcal{L}_C is correct. Consider any z in the domain of f . Since Z is empty at termination, z must be accepted by some conjunction in \mathcal{L}_C . Let (C, ℓ_i) be the first entry such that $C(z) = 1$, so $z \in Z$ during the iteration of the inner loop when this entry was added. Since in this iteration we have $G(A \times B) = C^{-1}(1)$ and $z \in C^{-1}(1)$, there is some $(x, y) \in A \times B$ with $G(x, y) = z$. Since $A \times B \subseteq R_i$ we have $(x, y) \in R_i$. Since $A \times B \subseteq X \times Y$, we have $(x, y) \in (X \times Y) \cap G^{-1}(Z)$ and thus (x, y) cannot be in R_j for any $j < i$ since $R_j \cap (X \times Y) \cap G^{-1}(Z) = \emptyset$ by the invariant. In summary, R_i is the first rectangle in \mathcal{L}_R that contains (x, y) . By correctness of \mathcal{L}_R on $\mathcal{X} \times \mathcal{Y}$, we have $\oplus(\ell_i) = (f \circ g^n)(x, y) = f(G(x, y)) = f(z)$. Thus \mathcal{L}_C also correctly outputs $\oplus(\ell_i)$ on input z .

The “moreover” part is straightforward to verify: the levels assigned to conjunctions in \mathcal{L}_C come from the levels assigned to rectangles in \mathcal{L}_R (namely $\{0, \dots, q\}$), in the same order (which is non-increasing). \square

6 Query-to-communication lifting for $\mathbf{P}_{\parallel}^{\text{NP}[q]}$

Restatement of Theorem 4.(ii). For every partial function $f: \{0, 1\}^n \rightarrow \{0, 1\}$ and every constant q ,

$$\mathbf{P}_{\parallel}^{\text{NP}[q]\text{cc}}(f \circ g^n) = \mathbf{P}_{\parallel}^{\text{NP}[q]\text{dt}}(f) \cdot \Theta(\log n)$$

where $g: [m] \times \{0, 1\}^m \rightarrow \{0, 1\}$ is the index gadget defined by $g(x, y) = y_x$ with $m := n^{20}$.

For the big- O direction, the deterministic phase of a $\mathbb{P}_{\parallel}^{\text{NP}(q)\text{dt}}$ decision tree can be simulated by a protocol that communicates $\log m + 1 = O(\log n)$ bits to evaluate $g(x_i, y_i)$ whenever the decision tree queries the i^{th} bit of the input to f . The NP^{dt} oracle queries can also be converted to NP^{cc} oracle queries with $O(\log n)$ factor overhead in their contribution to the cost (as was the case for [Theorem 4.\(i\)](#)). In the rest of this section we prove the big- Ω direction. By [Lemma 16](#) it suffices to show

$$\mathbb{P}^{\text{NP}(q)[1]\text{cc}}(f \circ g^n) = \mathbb{P}^{\text{NP}(q)[1]\text{dt}}(f) \cdot \Omega(\log n).$$

6.1 Technical preliminaries

The high-level idea is to convert a $\mathbb{P}^{\text{NP}(q)[1]\text{cc}}$ protocol for $f \circ g^n$ into a $\mathbb{P}^{\text{NP}(q)[1]\text{dt}}$ decision tree for f by using the \mathbb{P} lifting theorem of Raz and McKenzie [[RM99](#), [GPW18a](#)] to handle the deterministic phase, followed by our $\text{NP}(q)$ lifting theorem to handle the single $\text{NP}(q)$ oracle query. To get these components to mesh, we need to review some technical concepts from the deterministic lifting theorem.

We say that $A \subseteq [m]^n$ is δ -thick if A is nonempty and for every $i \in [n]$ and every $\alpha_1 \cdots \alpha_{i-1} \alpha_{i+1} \cdots \alpha_n \in [m]^{[n] \setminus \{i\}}$, if there is at least one value of $\alpha_i \in [m]$ for which the combined tuple α is in A , then there are at least m^δ many such values of α_i . Recall that for $A \subseteq [m]^n$, $B \subseteq (\{0, 1\}^m)^n$, and $I \subseteq [n]$, $A_I := \{x_I : x \in A\}$ and $B_I := \{y_I : y \in B\}$ are the *projections* onto the coordinates of I . Also recall that the *deficiency* of B is $\mathbf{D}_\infty(B) := mn - \log |B|$. For a partial assignment $\rho \in \{0, 1, *\}^n$, let $\text{free}(\rho) := \rho^{-1}(*) \subseteq [n]$ denote its free coordinates and $\text{fixed}(\rho) := [n] \setminus \text{free}(\rho)$ denote its fixed coordinates. The rectangle $A \times B$ is called ρ -consistent if every assignment in $g^n(A \times B)$ is consistent with ρ .

The deterministic lifting theorem of [[RM99](#), [GPW18a](#)] proves the following result.

Lemma 24. *For every deterministic communication protocol Π_{det} of cost k (with the same input domain as g^n where g is the index gadget with $m := n^{20}$), there exists a deterministic decision tree T_{det} of cost $\leq 40k/\log m$ (with input domain $\{0, 1\}^n$) such that the following holds: For every leaf of T_{det} , letting $\rho \in \{0, 1, *\}^n$ be the partial assignment recording the results of the queries on the path to this leaf (so $|\text{fixed}(\rho)| \leq 40k/\log m$), there exists a rectangle $A \times B \subseteq [m]^n \times (\{0, 1\}^m)^n$ such that:*

- $A \times B$ is contained within one of the leaf rectangles of Π_{det} .
- $A \times B$ is ρ -consistent.
- $A_{\text{free}(\rho)}$ is 0.85-thick.
- $\mathbf{D}_\infty(B_{\text{free}(\rho)}) \leq n^2$ (relative to $(\{0, 1\}^m)^{\text{free}(\rho)}$).

The idea is to apply [Lemma 24](#) to the deterministic phase of the $\mathbb{P}^{\text{NP}(q)[1]\text{cc}}$ protocol, and then apply [Lemma 23](#) to each of the leaves of the resulting deterministic decision tree. In order to do so, we need the following claim, which observes that 0.85-thickness implies 0.85-density.

Claim 25. *If $X \subseteq [m]^n$ is δ -thick then X is δ -dense.*

Proof. Let \mathbf{x} be the uniform random variable over X . The basic intuition is that thickness tells us that for any single coordinate i , the probability that \mathbf{x}_i takes on a particular value, given that all of the other coordinates are fixed, is at most $m^{-\delta}$. This implies that the probability that \mathbf{x}_i takes on a particular value, conditioned on setting any subset of the other coordinates to some values, is still bounded by $m^{-\delta}$. Therefore, we can simply apply the chain rule to prove that thickness implies density.

Assume for notational simplicity that I is the first k coordinates, so $I = \{1, \dots, k\}$. Let $\alpha = \alpha_1 \cdots \alpha_k \in [m]^I$ be an assignment to the coordinates of I . By the chain rule,

$$\Pr[\mathbf{x}_I = \alpha_1 \cdots \alpha_k] = \Pr[\mathbf{x}_1 = \alpha_1] \cdot \Pr[\mathbf{x}_2 = \alpha_2 \mid \mathbf{x}_1 = \alpha_1] \cdot \dots \cdot \Pr[\mathbf{x}_k = \alpha_k \mid \mathbf{x}_{I \setminus \{k\}} = \alpha_1 \cdots \alpha_{k-1}]$$

and each conditioning is valid, assuming $\Pr[\mathbf{x}_I = \alpha] > 0$.

Thickness of X tells us that for any coordinate $i \in I$, any $\alpha_i \in [m]$, and any $\beta \in [m]^{[n] \setminus \{i\}}$, we have $\Pr[\mathbf{x}_i = \alpha_i \mid \mathbf{x}_{[n] \setminus \{i\}} = \beta] \leq m^{-\delta}$ if the conditioning is valid. Thus for any $I' \subseteq [n] \setminus \{i\}$, and any assignment β' to the coordinates of I' , $\Pr[\mathbf{x}_i = \alpha_i \mid \mathbf{x}_{I'} = \beta'] \leq m^{-\delta}$ if the conditioning is valid. Therefore, each term in the chain rule expansion is at most $m^{-\delta}$. This gives $\Pr[\mathbf{x}_I = \alpha] \leq (m^{-\delta})^k = 2^{-\delta|I| \log m}$, which implies $\mathbf{H}_\infty(\mathbf{x}_I) \geq \delta|I| \log m$ since α was arbitrary. Since this holds for all nonempty subsets $I \subseteq [n]$, this means that X is δ -dense. \square

6.2 The simulation

Let Π be a $\mathsf{P}^{\mathsf{NP}(q)[1]^{\text{cc}}}$ protocol for $f \circ g^n$ with cost k . We construct a $\mathsf{P}^{\mathsf{NP}(q)[1]^{\text{dt}}}$ decision tree T for f with cost $O(k/\log n)$.

If $k = \Omega(n \log n)$, we can construct a deterministic decision tree for f with cost n that simply queries every bit in the input. This is a $\mathsf{P}^{\mathsf{NP}(q)[1]^{\text{dt}}}$ decision tree for f (with a trivial $\mathsf{NP}(q)^{\text{dt}}$ phase) with cost $O(k/\log n)$. In the following, assume that $k = o(n \log n)$.

Let Π_{det} be the deterministic phase of Π . Each leaf v of Π_{det} has an associated rectangle R_v and $\mathsf{NP}(q)^{\text{cc}}$ protocol Π_v of cost $\leq k$ that computes either $f \circ g^n$ or its complement on inputs in R_v . Applying [Lemma 24](#) to Π_{det} , which has communication cost $\leq k$, yields a deterministic decision tree T_{det} having query cost $\leq 40k/\log m = O(k/\log n)$, which will form the deterministic phase of T . Henceforth consider any leaf of T_{det} , say corresponding to partial assignment ρ , and from [Lemma 24](#) let v be the associated leaf of Π_{det} and $A \times B \subseteq R_v$ be the associated rectangle. Assume without loss of generality that Π_v computes $f \circ g^n$ on R_v and hence on $A \times B$ (a symmetric argument handles the case where Π_v computes $\bar{f} \circ g^n$ on R_v). Abbreviate $\text{free}(\rho)$ as J , and let $n' := |J|$. Since $|\text{fixed}(\rho)| \leq 40k/\log m \leq n/2$ (as we are assuming $k = o(n \log n)$), we have $n' \geq n/2$.

Since $A \times B$ is ρ -consistent, if we modify Π_v by intersecting each rectangle with $A \times B$ and then projecting to J , this yields an $\mathsf{NP}(q)^{\text{cc}}$ protocol that correctly computes $f_\rho \circ g^J$ on $A_J \times B_J$, where $f_\rho: \{0, 1\}^J \rightarrow \{0, 1\}$ is the restriction of f to ρ . After converting this protocol to a q -alternating rectangle decision list of cost $O(k)$ by [Lemma 13](#), we may apply [Lemma 23](#) by substituting f_ρ for f , n' for n , A_J for \mathcal{X} (since A_J is 0.85-thick), and B_J for \mathcal{Y} (since $\mathbf{D}_\infty(B_J) \leq n^2 \leq (n')^3$),¹ to get a q -alternating conjunction decision list for f_ρ of cost $O(k/\log n') = O(k/\log n)$. By [Lemma 13](#) again, there is an $\mathsf{NP}(q)^{\text{dt}}$ decision tree T_v with cost $O(k/\log n)$ that correctly computes f_ρ . Therefore, we can complete T by having our arbitrary leaf of T_{det} use T_v as its oracle query and output the same answer. Thus T computes f and has total cost $O(k/\log n)$.

Acknowledgments

Toniann Pitassi and Morgan Shirley were supported by NSERC. Thomas Watson was supported by NSF grant CCF-1657377.

References

- [BBJ⁺89] Alberto Bertoni, Danilo Bruschi, Deborah Joseph, Meera Sitharam, and Paul Young. Generalized Boolean hierarchies and Boolean hierarchies over RP. In *Proceedings of the 7th International Conference on Fundamentals of Computation Theory (FCT)*, pages 35–46. Springer, 1989.
- [Bei91] Richard Beigel. Bounded queries to SAT and the Boolean hierarchy. *Theoretical Computer Science*, 84(2):199–223, 1991.
- [BFS86] 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.
- [CH86] Jin-Yi Cai and Lane Hemachandra. The Boolean hierarchy: Hardware over NP. In *Proceedings of the 1st Structure in Complexity Theory Conference (STRUCTURES)*, pages 105–124. Springer, 1986.
- [CLV18] Arkadev Chattopadhyay, Shachar Lovett, and Marc Vinyals. Equality alone does not simulate randomness. Technical Report TR18-206, Electronic Colloquium on Computational Complexity (ECCC), 2018.
- [GKPW17] Mika Göös, Prithish Kamath, Toniann Pitassi, and Thomas Watson. Query-to-communication lifting for P^{NP} . In *Proceedings of the 32nd Computational Complexity Conference (CCC)*, pages 12:1–12:16. Schloss Dagstuhl, 2017.

¹A subtlety is that [Lemma 23](#) assumes $m := (n')^{20}$ whereas here we have $m := n^{20}$, but this tiny perturbation in parameters does not affect the result.

- [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.
- [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.
- [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.
- [GPW18a] Mika Göös, Toniann Pitassi, and Thomas Watson. Deterministic communication vs. partition number. *SIAM Journal on Computing*, 47(6):2435–2450, 2018.
- [GPW18b] Mika Göös, Toniann Pitassi, and Thomas Watson. The landscape of communication complexity classes. *Computational Complexity*, 27(2):245–304, 2018.
- [HR88] Bernd Halstenberg and Rüdiger Reischuk. Relations between communication complexity classes. In *Proceedings of the 3rd Structure in Complexity Theory Conference (STRUCTURES)*, pages 19–28. IEEE, 1988.
- [HR90] Bernd Halstenberg and Rüdiger Reischuk. Relations between communication complexity classes. *Journal of Computer and System Sciences*, 41(3):402–429, 1990.
- [HR93] Bernd Halstenberg and Rüdiger Reischuk. Different modes of communication. *SIAM Journal on Computing*, 22(5):913–934, 1993.
- [Juk12] Stasys Jukna. *Boolean Function Complexity: Advances and Frontiers*, volume 27 of *Algorithms and Combinatorics*. Springer, 2012.
- [Kad88] Jim Kadin. The polynomial time hierarchy collapses if the Boolean hierarchy collapses. *SIAM Journal on Computing*, 17(6):1263–1282, 1988.
- [KN97] Eyal Kushilevitz and Noam Nisan. *Communication Complexity*. Cambridge University Press, 1997.
- [KSW87] Johannes Köbler, Uwe Schöning, and Klaus Wagner. The difference and truth-table hierarchies for NP. *Theoretical Informatics and Applications*, 21(4):419–435, 1987.
- [Nis91] Noam Nisan. CREW PRAMs and decision trees. *SIAM Journal on Computing*, 20(6):999–1007, 1991.
- [Riv87] Ronald Rivest. Learning decision lists. *Machine Learning*, 2(3):229–246, 1987.
- [RM99] Ran Raz and Pierre McKenzie. Separation of the monotone NC hierarchy. *Combinatorica*, 19(3):403–435, 1999.
- [Wag88] Klaus Wagner. Bounded query computations. In *Proceedings of the 3rd Structure in Complexity Theory Conference (STRUCTURES)*, pages 260–277. IEEE, 1988.
- [Wat19] Thomas Watson. A $ZPP^{NP[1]}$ lifting theorem. In *Proceedings of the 36th International Symposium on Theoretical Aspects of Computer Science (STACS)*. Schloss Dagstuhl, 2019. To appear.
- [Wec85] Gerd Wechsung. On the Boolean closure of NP. In *Proceedings of the 5th International Conference on Fundamentals of Computation Theory (FCT)*, pages 485–493. Springer, 1985.
- [Yao79] Andrew Yao. Some complexity questions related to distributive computing. In *Proceedings of the 11th Symposium on Theory of Computing (STOC)*, pages 209–213. ACM, 1979.
- [ZH86] Stathis Zachos and Hans Heller. A decisive characterization of BPP. *Information and Control*, 69(1-3):125–135, 1986.

A Proofs from Section 2

In this appendix we prove the technical lemmas from Section 2, which provide alternative characterizations of the classes in the Nondeterministic and Randomized Boolean Hierarchies.

A.1 Decision list vs. parity

Restatement of Lemma 8. *For $\mathcal{C} \in \{\text{NP}, \text{RP}\}$, if the definitions of $\mathcal{C}(q)^{\text{cc}}$ and $\mathcal{C}(q)^{\text{dt}}$ are changed to use \oplus_q in place of Δ_q , it only affects the complexity measures $\mathcal{C}(q)^{\text{cc}}(F)$ and $\mathcal{C}(q)^{\text{dt}}(f)$ by a constant factor (depending on q).*

Wagner [Wag88] showed that these alternative characterizations are equivalent for the classical Nondeterministic Boolean Hierarchy, and Halstenberg and Reischuk [HR90] observed the same (up to a constant factor) in communication complexity. This latter proof uses only the property that NP^{cc} is closed under intersection and union; that is, if $\text{NP}^{\text{cc}}(F_1), \text{NP}^{\text{cc}}(F_2) \leq k$, then $\text{NP}^{\text{cc}}(F_1 \wedge F_2)$ and $\text{NP}^{\text{cc}}(F_1 \vee F_2)$ are both $O(k)$. We observe that since this property also holds for RP^{cc} , NP^{dt} , and RP^{dt} , their proof works for these models of computation as well. In fact, in all of these models, the cost of the intersection or union of i cost- k computations is at most ik .

Proof of Lemma 8. We prove this using the language of communication complexity. The query complexity proof is completely analogous.

Given a protocol $\Pi = \Delta_q(\Pi_1, \dots, \Pi_q)$ of cost k , we can transform it into an equivalent protocol $\Pi' = \oplus_q(\Pi'_1, \dots, \Pi'_q)$ by having Π'_i compute “does there exist a $j \geq i$ for which Π_j outputs 1?” This works because if i is the greatest index such that Π_i outputs 1, then Π outputs $\oplus(i)$ and Π' also outputs $\oplus(i)$ since there are exactly i many values of j such that Π'_j outputs 1, namely $1, \dots, i$. For the implementation of Π'_i : If $\mathcal{C} = \text{NP}$ then a witness for Π'_i can specify a $j \geq i$ together with a witness for Π_j , incurring a cost of $O(k)$; thus the cost of Π' is $O(qk) = O(k)$. If $\mathcal{C} = \text{RP}$ then Π'_i can run Π_j for every $j \geq i$, each amplified to have error probability $\leq 1/2q$, and see whether at least one of them outputs 1, incurring a cost of $O(k \log q)$; thus the cost of Π' is $O(qk \log q) = O(k)$.

Conversely, given a protocol $\Pi = \oplus_q(\Pi_1, \dots, \Pi_q)$ of cost k , we can transform it into an equivalent protocol $\Pi' = \Delta_q(\Pi'_1, \dots, \Pi'_q)$ by having Π'_i compute “are there at least i many values of j for which Π_j outputs 1?” This works because if there are exactly i values of j for which Π_j outputs 1, then Π outputs $\oplus(i)$ and Π' also outputs $\oplus(i)$ since i is the greatest index such that Π'_i outputs 1. For the implementation of Π'_i : If $\mathcal{C} = \text{NP}$ then a witness for Π'_i can specify a set of i values of j together with a witness for each of those Π_j , incurring a cost of $O(k)$; thus the cost of Π' is $O(qk) = O(k)$. If $\mathcal{C} = \text{RP}$ then Π'_i can run every Π_j , each amplified to have error probability $\leq 1/2q$, and see whether at least i of them output 1, incurring a cost of $O(k \log q)$; thus the cost of Π' is $O(qk \log q) = O(k)$. \square

A.2 Boolean Hierarchy vs. decision list

Restatement of Lemma 13. *$\text{DL}(q)^{\text{cc}}(F) = \Theta(\text{NP}(q)^{\text{cc}}(F))$ and $\text{DL}(q)^{\text{dt}}(f) = \Theta(\text{NP}(q)^{\text{dt}}(f))$ for every constant q . Thus, $\text{DL}(q)^{\text{cc}} = \text{NP}(q)^{\text{cc}}$ and $\text{DL}(q)^{\text{dt}} = \text{NP}(q)^{\text{dt}}$ for partial functions.*

Proof. To see that $\text{DL}(q)^{\text{cc}}(F) \leq \text{NP}(q)^{\text{cc}}(F)$, consider any $\text{NP}(q)^{\text{cc}}$ protocol for F with cost k , say $\Pi = (\mathcal{R}_1, \dots, \mathcal{R}_q)$ where each \mathcal{R}_i is a nonempty set of rectangles. To form a q -alternating rectangle decision list, let level q be the rectangles of \mathcal{R}_q in any order, then level $q-1$ be the rectangles of \mathcal{R}_{q-1} in any order, and so on, and finally let level 0 be the rectangle containing all inputs. This has the same output as Π , so it correctly computes F . Assuming \mathcal{R}_i has cost k_i , the length of the list is $\sum_i |\mathcal{R}_i| + 1 \leq \prod_i (|\mathcal{R}_i| + 1) \leq \prod_i 2^{k_i} = 2^k$, so the cost is at most k .

To see that $\text{NP}(q)^{\text{cc}}(F) = O(\text{DL}(q)^{\text{cc}}(F))$, consider any q -alternating rectangle decision list \mathcal{L}_R for F with cost k . To form an $\text{NP}(q)^{\text{cc}}$ protocol $\Pi = (\mathcal{R}_1, \dots, \mathcal{R}_q)$, for each i let \mathcal{R}_i be the set of rectangles at level i in \mathcal{L}_R . This has the same output as \mathcal{L}_R , so it correctly computes F . Since $|\mathcal{R}_i| \leq 2^k - 1$ for each i , the cost of Π is at most $\sum_i k = qk$.

To see that $\text{DL}(q)^{\text{dt}}(f) \leq \text{NP}(q)^{\text{dt}}(f)$, consider any $\text{NP}(q)^{\text{dt}}$ decision tree for f with cost k , say $T = (\Phi_1, \dots, \Phi_q)$ where each Φ_i is a DNF. To form a q -alternating conjunction decision list, let level q be the conjunctions of Φ_q in any order, then level $q-1$ be the conjunctions of Φ_{q-1} in any order, and so on, and

finally let level 0 be the conjunction that accepts all inputs. This has the same output as T , so it correctly computes f . Since every Φ_i has maximum width at most k , the list also has cost at most k .

To see that $\text{NP}(q)^{\text{dt}}(f) = O(\text{DL}(q)^{\text{dt}}(f))$, consider any q -alternating conjunction decision list \mathcal{L}_C for f with cost k . To form an $\text{NP}(q)^{\text{dt}}$ decision tree $T = (\Phi_1, \dots, \Phi_q)$, for each i let Φ_i be the disjunction of all conjunctions at level i in \mathcal{L}_C . This has the same output as \mathcal{L}_C , so it correctly computes f . Since each Φ_i has maximum width at most k , the cost of T is at most $\sum_i k = qk$. \square

A.3 Leaves with parallel queries

Restatement of Lemma 16. *For $\mathcal{C} \in \{\text{NP}, \text{RP}\}$, we have $\text{P}_{\parallel}^{\mathcal{C}[q]^{\text{cc}}}(F) = \Theta(\text{P}^{\mathcal{C}(q)[1]^{\text{cc}}}(F))$ and $\text{P}_{\parallel}^{\mathcal{C}[q]^{\text{dt}}}(f) = \Theta(\text{P}^{\mathcal{C}(q)[1]^{\text{dt}}}(f))$ for every constant q .*

Proof. It is trivial that $\text{P}_{\parallel}^{\mathcal{C}[q]^{\text{cc}}}(F) \leq \text{P}^{\mathcal{C}(q)[1]^{\text{cc}}}(F)$ and $\text{P}_{\parallel}^{\mathcal{C}[q]^{\text{dt}}}(f) \leq \text{P}^{\mathcal{C}(q)[1]^{\text{dt}}}(f)$ since one “ $\mathcal{C}(q)$ oracle query” can be expressed with q nonadaptive “ \mathcal{C} oracle queries”. For the other direction, we just show the argument for NP^{cc} , but essentially the same argument works for RP^{cc} , NP^{dt} , and RP^{dt} .

Consider a $\text{P}_{\parallel}^{\text{NP}(q)^{\text{cc}}}$ protocol Π for F of cost k . We convert it to a $\text{P}^{\text{NP}(q)[1]^{\text{cc}}}$ protocol for F of cost $O(k)$. The deterministic phase Π_{det} stays the same. Henceforth fix any leaf of Π_{det} and its associated output function $\text{out}: \{0, 1\}^q \rightarrow \{0, 1\}$ and q nonempty sets of rectangles $\mathcal{R}_1, \dots, \mathcal{R}_q$. Also fix any input (x, y) that reaches this leaf, and define $w \in \{0, 1\}^q$ by $w_i = \mathcal{R}_i(x, y)$ for each $i \in [q]$, so that $\text{out}(w) = \Pi(x, y) = F(x, y)$.

For $u, v \in \{0, 1\}^q$ we say $u \leq v$ iff $u_i \leq v_i$ for each index $i \in [q]$. Consider the q -dimensional *Hamming cube DAG*: this is the graph where the set of nodes is $\{0, 1\}^q$ and there is a directed edge from u to v iff $u \leq v$ and u and v only differ by one index. Each node v has a corresponding output value $\text{out}(v)$, and on any directed path $v^1 \rightarrow v^2 \rightarrow \dots \rightarrow v^j$, we say there is a *mind-change* on the path at node v^i iff $\text{out}(v^i) \neq \text{out}(v^{i-1})$. For any directed path in the Hamming cube DAG from 0^q to v , the number of mind-changes on the path is even if $\text{out}(v) = \text{out}(0^q)$ and odd if $\text{out}(v) \neq \text{out}(0^q)$. Thus, letting $\text{mmc}(v)$ be the maximum number of mind-changes on any directed path from 0^q to v , we have $\oplus(\text{mmc}(v))$ is the indicator for $\text{out}(v) \neq \text{out}(0^q)$.

Since $\text{out}(0^q)$ is fixed at the leaf, it is sufficient for Alice and Bob to determine $\oplus(\text{mmc}(w))$ in order to compute $\text{out}(w) = F(x, y)$.

Claim 26. *For fixed $i \in [q]$, the cost of an NP^{cc} protocol that determines “is $\text{mmc}(w) \geq i$?” (given input (x, y) that reaches the leaf) is at most $q + k$.*

Before we prove **Claim 26**, we use it to finish the proof of **Lemma 16**. At the leaf, for each $i \in [q]$ we form an NP^{cc} protocol of cost $\leq q + k$ for determining if $\text{mmc}(w) \geq i$, and we use these to form a single $\text{NP}(q)^{\text{cc}}$ oracle query. We output the same answer as the oracle if $\text{out}(0^q) = 0$, and the opposite answer if $\text{out}(0^q) = 1$; this is correct because $\text{mmc}(w)$ is the maximum value of i for which the i^{th} NP^{cc} protocol would output 1, and thus the oracle query returns $\oplus(\text{mmc}(w))$, which is the indicator for $\text{out}(w) \neq \text{out}(0^q)$. The $\text{NP}(q)^{\text{cc}}$ oracle query has cost $\leq q(q + k) = O(k)$. Thus the overall $\text{P}^{\text{NP}(q)[1]^{\text{cc}}}$ protocol for F has cost $O(k)$, and this concludes the proof of **Lemma 16**. \square

Proof of Claim 26. The idea is to find some (possibly not proper) prefix of a directed path from 0^q to w in the Hamming cube DAG, where this prefix has i mind-changes. If we can find such a prefix, it confirms that some 0^q -to- w directed path has at least i mind-changes, and therefore $\text{mmc}(w) \geq i$.

The witness itself is some $v \in \{0, 1\}^q$, along with witnesses of $\mathcal{R}_j(x, y) = 1$ for every j where $v_j = 1$. This can be represented by simply giving a rectangle R' such that $(x, y) \in R'$ and R' is the intersection of (Hamming weight of v many) rectangles, one from each \mathcal{R}_j with $v_j = 1$. The number of possibilities for R' is at most $\prod_j |\mathcal{R}_j| \leq 2^k$, so R' contributes $\leq k$ to the cost of the NP^{cc} protocol. Alice and Bob can verify the witness by checking that there exists a 0^q -to- v path with i mind-changes (this does not depend on the input) and that $(x, y) \in R'$ and thus $v \leq w$. \square