Electronic Colloquium on Computational Complexity, Report No. 166 (2023)

Patterned non-determinism in communication complexity*

Dmytro Gavinsky†‡

November 5, 2023

Abstract

We define and study the model of *patterned non-determinism* in bipartite communication complexity, denoted by $PNP^{X\leftrightarrow Y}$. It generalises the known models $UP^{X\leftrightarrow Y}$ and $FewP^{X\leftrightarrow Y}$ through relaxing the constraints on the witnessing structure of the underlying $NP^{X\leftrightarrow Y}$-protocol.

It is shown that for the case of total functions $PNP^{X\leftrightarrow Y}$ equals $P^{X\leftrightarrow Y}$ (similarly to $UP^{X\leftrightarrow Y}$ and $FewP^{X\leftrightarrow Y}$). Moreover, the corresponding *exhaustive witness-searching* problem – determining the *full set* of witnesses that lead to the acceptance of a given input pair – also has an efficient deterministic protocol.

Structurally, the possibility of efficient exhaustive $PNP^{X\leftrightarrow Y}$-search summarises the above results and can be stated like this: if $f_1,\dots,f_m$ are bipartite total Boolean functions with efficient deterministic protocols, then for every input (x,y) the set $\left\{i\middle|f_i(x,y)=\top\right\}$ can be found by a deterministic protocol of cost *poly-logarithmic* in n and the total number of such sets for these f_i's.

Finally, the possibility of efficient exhaustive $PNP^{X\leftrightarrow Y}$-search is used to analyse certain *three-party communication* regime (under the "number in hand" input partition): The corresponding three-party model is shown to be as strong qualitatively as the weakest among its two-party amplifications obtained by allowing free communication between a pair of players.

1 Introduction

Let $f(x,y):\{0,1\}^n\times\{0,1\}^n\to\{\top,\bot\}$ be a total bipartite communication problem with an efficient $NP^{X\leftrightarrow Y}$-protocol Π, that is, the total number of bits sent by $\Pi(x,y)$ is in poly-log(n). It was shown by Yannakakis [Yan91] that if for every $(x,y)\in f^{-1}(\top)$ there is exactly one Π-witness, then $f\in P^{X\leftrightarrow Y}$. Later Karchmer, Newman, Saks and Wigderson [KNSW94] strengthened the result by drawing the same conclusion from the weaker assumption that the number of Π-witnesses per input pair was at most poly-log(n). The corresponding communication complexity classes – that is, the families of functions for which there are efficient protocols – are denoted by $UP^{X\leftrightarrow Y}$ and $FewP^{X\leftrightarrow Y}$ and the above results can be stated as $UP^{X\leftrightarrow Y}=FewP^{X\leftrightarrow Y}=P^{X\leftrightarrow Y}$.[1]

Consider the following generalisation. Let $\mathcal{W}$ be the family of all possible Π's witnesses and call $\gamma\subseteq\mathcal{W}$ a *pattern* (for the protocol Π) if for some $(x,y)\in f^{-1}(\top)$ the set of witnesses that cause Π's

*This is a preliminary version...

†In 2022 the author has changed the English spelling of his first name from the previous russian-odoured form *"Dmitry"* to the Ukrainian *"Dmytro"*.

‡Institute of Mathematics of the Czech Academy of Sciences, Žitna 25, Praha 1, Czech Republic.

Partially funded by the grant 19-27871X of GA ČR and by RVO: 67985840. Part of this work was done while visiting the Centre for Quantum Technologies at the National University of Singapore.

[1] To denote communication complexity *classes*, as well as the corresponding *models*, we will usually add the superscript "$X\leftrightarrow Y$" to the common notation for the corresponding computational complexity class.

ISSN 1433-8092

acceptance of (x, y) equals γ. If the total number of Π's patterns is at most $2^{k(n)}$ for $k(n) \in \text{poly-log}(n)$, then we say that Π is an efficient protocol in the model of *patterned non-determinism*, and the Boolean function f that Π computes belongs to the corresponding communication complexity class $PNP^{X \leftrightarrow Y}$ (obviously, $UP^{X \leftrightarrow Y} \subseteq FewP^{X \leftrightarrow Y} \subseteq PNP^{X \leftrightarrow Y}$). We will see that $PNP^{X \leftrightarrow Y} = P^{X \leftrightarrow Y}$.

Next we consider the communication complexity of the *exhaustive witness-search problem* corresponding to $PNP^{X \leftrightarrow Y}$-protocols (or simply *exhaustive $PNP^{X \leftrightarrow Y}$-search*): by this we will mean determining the *full set* of witnesses that lead to the acceptance of a given input pair by the given $PNP^{X \leftrightarrow Y}$-protocol (which is, in particular, an $NP^{X \leftrightarrow Y}$-protocol).

On the one hand, we will see that from the equality of a certain subclass $\mathscr{C} \subseteq NP^{X \leftrightarrow Y}$ to $P^{X \leftrightarrow Y}$ even the possibility of efficiently finding *any protocol-compatible witness* doesn't follow in general (leave alone determining the complete set of valid witnesses).[2] Nevertheless, efficient deterministic witness-searching for $PNP^{X \leftrightarrow Y}$-protocols will be presented. That is, for an efficient $NP^{X \leftrightarrow Y}$-protocol Π with at most $2^{\text{poly-log}(n)}$ patterns there exists a $P^{X \leftrightarrow Y}$-protocol Π_{search} of cost at most $\text{poly-log}(n)$ that finds – for every input pair (x, y) that Π accepts – *the exact set of witnesses* that lead to Π's acceptance of (x, y).[3]

More formally, if $f_1(x, y), \ldots, f_m(x, y) : \{0,1\}^n \times \{0,1\}^n \to \{\top, \bot\}$ are such that every f_i has an efficient deterministic protocol and the set $\left\{\left\{i \middle| f_i(x, y) = \top\right\}\right\}_{x,y}$ is of size at most quasi-polynomial in n, then the exhaustive-search function $F(x, y) \stackrel{\text{def}}{=} \left\{i \middle| f_i(x, y) = \top\right\}$ is in $P^{X \leftrightarrow Y}$ (i.e., there is an efficient deterministic protocol). Note that this statement generalises the equality $PNP^{X \leftrightarrow Y} = P^{X \leftrightarrow Y}$.

Finally, let $g(x, y, z) : \{0,1\}^n \times \{0,1\}^n \times \{0,1\}^n \to \{0,1\}$ be a tripartite total function. Consider the following scenarios of deterministically computing g:[4]

- Denote by $[(X \leftrightarrow Y) \to Z]$ the regime where Alice receives x, Bob receives y, Charlie receives z, Alice and Bob interact in order to produce a message that is sent to Charlie, who must answer upon receiving it (alternatively, this setting can be viewed as having "broadcasting" interaction between Alice and Bob, that is, letting Charlie see its transcript).
- Denote by $[(X, Y) \to Z]$ the regime where Alice receives (x, y), Charlie receives z, Alice sends a message to Charlie, who must answer upon receiving it.
- Denote by $[(X, Z) \leftrightarrow (Y, Z)]$ the regime where Alice receives (x, z), Bob receives (y, z), they interact until Bob produces the answer.

For brevity we will say that $g(x, y, z)$ *is efficiently computable in* $[(X, Z) \leftrightarrow (Y, Z)] \cap [(X, Y) \to Z]$ if it has efficient protocols in both $[(X, Z) \leftrightarrow (Y, Z)]$ and $[(X, Y) \to Z]$.

Note that computing a tripartite function in $[(X \leftrightarrow Y) \to Z]$ is at least as hard as computing it in $[(X, Z) \leftrightarrow (Y, Z)] \cap [(X, Y) \to Z]$: Assume the existence of an efficient three-player protocol Π in the setting $[(X \leftrightarrow Y) \to Z]$, then the existence of an efficient protocol in $[(X, Y) \to Z]$ follows by "merging" Alice and Bob in Π (i.e., letting them communicate for free) and an efficient protocol in $[(X, Z) \leftrightarrow (Y, Z)]$ can be obtained by "merging" Bob and Charlie (the resulting protocol only uses

[2] Even though $\mathscr{C} \subseteq P^{X \leftrightarrow Y}$ trivially implies the existence of *some* efficiently verifiable witness for every answer to a problem $f(x, y) \in \mathscr{C}$, it can be the case that witnessing in accordance with the same $NP^{X \leftrightarrow Y}$-protocol that establishes the membership of f in $\mathscr{C}$ is not feasible: We will see that there are subsets $\mathcal{A}, \mathcal{B} \subseteq \{0,1\}^n$ such that the (total) instance of the *set intersection problem* defined over $\mathcal{A} \times \mathcal{B}$ belongs to $NP^{X \leftrightarrow Y} \cap coNP^{X \leftrightarrow Y}$ – therefore to $P^{X \leftrightarrow Y}$ [AUY83] – while at the same time finding a presumably existing index $i \in [n]$ such that $x_i = y_i = 1$ is not only infeasible deterministically, but also hard for randomised protocols over the uniformly-random input from $\left\{(x, y) \in \mathcal{A} \times \mathcal{B} \middle| \exists\, i :\ x_i = y_i = 1\right\}$.

[3] Note that the set can, in general, be large, containing as many elements as there are different witnesses in Π.

[4] The intuition behind the notation used next is clear: "$[(X, Y) \to Z]$" means that a holder of X and Y sends a 1-way message to the holder of Z, "$[(X, Z) \leftrightarrow (Y, Z)]$" means that a holder of X and Z is allowed to interact in the 2-way regime with the holder of Y and Z and so on.

even more restricted setting that can be denoted by $[X \leftrightarrow (Y, Z)]$, where Alice doesn't receive Z as part of her input).

We will see that every total function $g(x, y, z)$ that is efficiently computable in $[(X, Z) \leftrightarrow (Y, Z)] \cap [(X, Y) \to Z]$ has an efficient protocol in $[(X \leftrightarrow Y) \to Z]$ as well: this will follow from the possibility of efficient exhaustive $PNP^{X \leftrightarrow Y}$-search. In particular, the model of *deterministic interactive Alice and Bob with listening Charlie*, $[(X \leftrightarrow Y) \to Z]$, is as strong qualitatively as the weakest among its two-party amplifications obtained by allowing free communication between a pair of players.

Related work

For both *non-deterministic* and *randomised* setting in the *unrestricted* interactive three-party case, Draisma, Kushilevitz and Weinreb [DKW11] have demonstrated an *exponential gap* between the communication complexity of a tripartite total function and the largest of its three bipartite complexities in the amplified models resulting from allowing free communication between a pair of players. The new three-party result (Section 4) can be viewed as complementary to [DKW11]: it shows that the gap is *at most polynomial* in the case of *deterministic* interactive Alice and Bob with *listening Charlie*. The case of three *deterministic* players with *unrestricted* interaction remains open.

2 Preliminaries and definitions

We will write $[n]$ to denote the set $\{1, \ldots, n\} \subset \mathbb{N}$. Let (a, b), $[a, b]$, $[a, b)$ and $(a, b]$ denote the corresponding open, closed and half-open intervals in $\mathbb{R}$. For a finite $S \subset \mathbb{N}$ we will write $S(i)$ to address the i'th element of S in natural ordering. For any set S we will denote by $\mathsf{pow}(S)$ the family of its subsets and by $\binom{S}{t}$ the family of size-t subsets. We will write $x \in\!\!\!\circledcirc S$ to say that x is a uniformly random element of S. Towards readability, we will allow both $\{\cdot|\cdot\}$ and $\{\cdot : \cdot\}$ to denote sets with conditions (preferring the former).

For $x \in \{0,1\}^n$ and $i \in [n]$, we will write x_i or $x(i)$ to address the i'th bit of x (preferring "x_i" unless it may cause ambiguity). Let $|x|$ denote the Hamming weight of x. At times we will implicitly assume (without causing ambiguity) the trivial isomorphism between the n-bit strings and the subsets of $[n]$: in particular, the notation $\binom{[n]}{k}$ will stand for $\{x \in \{0,1\}^n \,||x| = k\}$, and $x \cap y$ will address the set $\{i \in [n] | x_i = y_i = 1\}$.

Let $\bot$ and $\top$ denote, respectively, the false and the true values: sometimes we will use the Boolean domain $\{\bot, \top\}$ (instead of $\{0,1\}$) to emphasise the intuitive asymmetry between the two values (say, when the non-deterministic computation is distinguished from the co-non-deterministic one, or if there is a "clear logical flavour" inherent to the values).

By default the logarithms are base-2.

We will use $\leftarrowtail$ to denote the assignment operation (e.g., in algorithms).

2.1 Communication complexity

The study of communication complexity was initiated by Abelson [Abe78] in the regime of real-valued messages and adapted by Yao [Yao79] to the discrete regime that we are interested in. We refer the reader to [KN97] for a classical background on communication complexity in general, to [GPW18] for a great survey of the more recent structural developments and to [DKW11] for some insight into the multi-party communication complexity setting.

Unless stated otherwise, the communication problems considered in this work are *total functions* (the only exception will be *witness-search problems*).

We will add the superscript "$X \leftrightarrow Y$" to the common notation for a computational complexity class to denote the corresponding communication complexity class (e.g., $P^{X \leftrightarrow Y}$ or $NP^{X \leftrightarrow Y}$). The resulting symbol will be used in three ways: to address the class itself; to address the corresponding communication model; to denote the complexity of a communication problem in that model (e.g., $P^{X \leftrightarrow Y}(f)$ is the deterministic communication complexity of f).

As the standard models $P^{X \leftrightarrow Y}$ and $NP^{X \leftrightarrow Y}$ are of core importance for this work, their definitions for the case of total functions are given next for the reader's convenience.

Definition 1 ($P^{X \leftrightarrow Y}$*, deterministic two-party communication*)**.** *For $n \in \mathbb{N}$, let the sets $|\mathcal{A}|, |\mathcal{B}|$ be such that $\max\{|\mathcal{A}|, |\mathcal{B}|\} \in (2^{n-1}, 2^n]$ and let $f : \mathcal{A} \times \mathcal{B} \to \{0,1\}$.*

Let Π be a deterministic protocol where

- *Alice receives x and Bob receives y;*
- *Alice and Bob interact;*
- *Bob produces the answer.*

If the transcript of $\Pi(x,y)$ contains at most $k(n)$ bits and the protocol computes $f(x,y)$, then we say that the $P^{X \leftrightarrow Y}$-complexity of f, denoted by $P^{X \leftrightarrow Y}(f)$, is at most $k(n)$.

We call a protocol efficient if its transcript contains at most poly-log(n) *bits and we say that a function is efficiently computable in $P^{X \leftrightarrow Y}$ if it has an efficient $P^{X \leftrightarrow Y}$-protocol. We denote by $P^{X \leftrightarrow Y}$ the class of total bipartite Boolean functions (or, alternatively, the languages of satisfying assignments to such functions, viewed as predicates) that are efficiently computable in $P^{X \leftrightarrow Y}$.*

Definition 2 ($NP^{X \leftrightarrow Y}$*, non-deterministic two-party communication*)**.** *Let R be a family of combinatorial rectangles in $\mathcal{A} \times \mathcal{B}$, $|R| \in (2^{k(n)-1}, 2^{k(n)}]$. Denote by Π_R the corresponding $NP^{X \leftrightarrow Y}$-protocol: it has complexity $k(n)$ and computes the predicate*

$$f_R(x,y) \stackrel{def}{=} \begin{cases} \top & \text{if } (x,y) \in r \text{ for some } r \in R; \\ \bot & \text{otherwise.} \end{cases}$$

For $n \in \mathbb{N}$, let the sets $|\mathcal{A}|, |\mathcal{B}|$ be such that $\max\{|\mathcal{A}|, |\mathcal{B}|\} \in (2^{n-1}, 2^n]$ and let $f : \mathcal{A} \times \mathcal{B} \to \{\top, \bot\}$. The $NP^{X \leftrightarrow Y}$-complexity of f, denoted by $NP^{X \leftrightarrow Y}(f)$, equals the minimal complexity of an $NP^{X \leftrightarrow Y}$-protocol that computes $f(x,y)$. We denote by $NP^{X \leftrightarrow Y}$ the class of total bipartite Boolean functions (or, alternatively, the languages of satisfying assignments to such functions, viewed as predicates) whose $NP^{X \leftrightarrow Y}$-complexity is at most poly-log(n).

3 Patterned non-determinism

While some of the definitions given next could be naturally generalised to the case of *partial* bipartite problems, we keep the notation simple by only considering the *total* case, which is of interest to us in this work. That is, the input space will have the product structure $\mathcal{A} \times \mathcal{B}$.

Definition 3 (*Accepting patterns of $NP^{X \leftrightarrow Y}$-protocols*)**.** *Let Π be an $NP^{X \leftrightarrow Y}$-protocol over input space $\mathcal{A} \times \mathcal{B}$ and let R_Π be the set of its rectangles.*

Call

$$\Gamma_\Pi \stackrel{def}{=} \left\{ \left\{ r \in R_\Pi \middle| (x,y) \in r \right\} \middle| (x,y) \in \mathcal{A} \times \mathcal{B} \right\}$$

the family of Π's accepting patterns.

Definition 4 ($PNP^{X\leftrightarrow Y}_{\square}$, *rectangle-patterned* $NP^{X\leftrightarrow Y}$)**.** *For $n \in \mathbb{N}$, let the sets $|\mathcal{A}|, |\mathcal{B}|$ be such that $\max\{|\mathcal{A}|, |\mathcal{B}|\} \in (2^{n-1}, 2^n]$ and let $f : \mathcal{A} \times \mathcal{B} \to \{\top, \bot\}$.*

Let Π be an $NP^{X\leftrightarrow Y}$-protocol (of any cost) that computes f such that the corresponding family of accepting patterns Γ_Π contains at most $2^{k(n)}$ elements, then we say that the $PNP^{X\leftrightarrow Y}_{\square}$-complexity of f, denoted by $PNP^{X\leftrightarrow Y}_{\square}(f)$, is at most $k(n)$.

We denote by $PNP^{X\leftrightarrow Y}_{\square}$ the class of total bipartite Boolean functions (or, alternatively, the languages of satisfying assignments to such functions, viewed as predicates) whose $PNP^{X\leftrightarrow Y}_{\square}$-complexity is at most $\text{poly-log}(n)$.

Note that the above definition does not require that the $NP^{X\leftrightarrow Y}$-protocol Π used to witness the $PNP^{X\leftrightarrow Y}_{\square}$-complexity of f is by itself efficient.[5]

The model $PNP^{X\leftrightarrow Y}_{\square}$ is a variation of previously studied $UP^{X\leftrightarrow Y}$, $FewP^{X\leftrightarrow Y}_t$ and $FewP^{X\leftrightarrow Y}$: they correspond to restricting Definition 4 by the condition that every $\gamma \in \Gamma_\Pi$ is of size at most 1, t or $\text{poly-log}(n)$, respectively. Trivially, $PNP^{X\leftrightarrow Y}_{\square}$ is a strengthening of those models (as long as $t \leq \text{poly-log}(n)$ in the case of $FewP^{X\leftrightarrow Y}_t$). On the other hand, Yannakakis [Yan91] proved that $UP^{X\leftrightarrow Y} = P^{X\leftrightarrow Y}$ and later Karchmer, Newman, Saks and Wigderson [KNSW94] strengthened it to $FewP^{X\leftrightarrow Y} = P^{X\leftrightarrow Y}$:[6]

Fact 1 ($FewP^{X\leftrightarrow Y}$ *vs.* $P^{X\leftrightarrow Y}$ *[KNSW94]*)**.** *For every total Boolean $f : \mathcal{A} \times \mathcal{B} \to \{\top, \bot\}$ and $t \in \mathbb{N}$,*

$$P^{X\leftrightarrow Y}(f) \in O\Big(t^2 \cdot FewP^{X\leftrightarrow Y}_t(f)^2\Big).$$

Accordingly, $FewP^{X\leftrightarrow Y} = P^{X\leftrightarrow Y}$.

In Section 3.1 we will address the question whether $PNP^{X\leftrightarrow Y}_{\square} = P^{X\leftrightarrow Y}$.

But there is a noteworthy intuitive difference between $PNP^{X\leftrightarrow Y}_{\square}$ and the classes $UP^{X\leftrightarrow Y}$, $FewP^{X\leftrightarrow Y}_t$ and $FewP^{X\leftrightarrow Y}$, namely *the robustness with respect to interactive verification* of the corresponding definitions. In Definitions 3 and 4 we treat individual rectangles of Π as the $NP^{X\leftrightarrow Y}$-witnesses, and it is natural to ask *what would happen to the defined models if, instead, we let the deterministic "verifier" be interactive*? That is, let Π' be an efficient deterministic protocol where Alice receives both $x \in \mathcal{A}$ and a witness $w \in \{0,1\}^{\text{poly-log}(n)}$, Bob receives $y \in \mathcal{B}$, then they interact and either accept or reject; say that such Π' computes the predicate $f_{\Pi'}(x_0, y_0)$ that gets the true value if and only if there exists w_0 such that $\Pi'((x_0, w_0), y_0)$ accepts.

Obviously, $f_{\Pi'} \in NP^{X\leftrightarrow Y}$. If it is additionally guaranteed that $\forall\,(x,y)\ \big|\{i | f_i(x,y) = \top\}\big| \leq 1$, $\leq t$ or $\leq \text{poly-log}(n)$, then, respectively, $f \in UP^{X\leftrightarrow Y}$, $f \in FewP^{X\leftrightarrow Y}_t$ or $f \in FewP^{X\leftrightarrow Y}$ – trivially, as follows from the respective definitions. The case of $PNP^{X\leftrightarrow Y}_{\square}$ is probably more interesting, as we see next.

For every $f \in NP^{X\leftrightarrow Y}$ we will assume a disjunctive decomposition $f(x,y) = \vee_{i=1}^m f_i(x,y)$, where every f_i represents the computation of the $NP^{X\leftrightarrow Y}$-protocol for the fixed witness value $w = i$ – that is, $\forall\, i :\ f_i \in P^{X\leftrightarrow Y}$.

Definition 5 (*Accepting patterns of disjunctions*)**.** *Let $f(x,y) = \vee_{i=1}^m f_i(x,y)$ be defined over $(x,y) \in \mathcal{A} \times \mathcal{B}$.*

[5] There are at most $|\Gamma_\Pi|$ rectangles that are not covered by other rectangles, so by dropping "meaningless" rectangles recursively, any Π can be transformed into an equivalent protocol of cost at most $\log(|\Gamma_\Pi|)$.

[6] Remember that all communication problems considered in this paper are total functions (for promise problems the equalities do not hold in general).

Call

$$\Gamma_f \stackrel{def}{=} \left\{ \left\{ i \middle| f_i(x,y) = \top \right\} \middle| (x,y) \in \mathcal{A} \times \mathcal{B} \right\}$$

the family of f's accepting patterns *with respect to the decomposition $\vee_i f_i(x,y)$ (often implicitly assumed).*

Definition 6 ($PNP^{X \leftrightarrow Y}$, *patterned* $NP^{X \leftrightarrow Y}$)**.** *For $n \in \mathbb{N}$, let the sets $|\mathcal{A}|, |\mathcal{B}|$ be such that $\max\{|\mathcal{A}|, |\mathcal{B}|\} \in (2^{n-1}, 2^n]$ and let $f : \mathcal{A} \times \mathcal{B} \to \{\top, \bot\}$.*

If f has a decomposition

$$f(x,y) \equiv \bigvee_{i=1}^{m} f_i(x,y),$$

such that the corresponding family of accepting patterns $\Gamma_f \subseteq \mathsf{pow}([m])$ contains at most $2^{k(n)}$ elements and $\forall\, i : \; P^{X \leftrightarrow Y}(f_i) \leq k(n)$, then we say that the $PNP^{X \leftrightarrow Y}$-complexity of f, denoted by $PNP^{X \leftrightarrow Y}(f)$, is at most $k(n)$. [7]

We denote by $PNP^{X \leftrightarrow Y}$ the class of total bipartite Boolean functions (or, alternatively, the languages of satisfying assignments to such functions, viewed as predicates) whose $PNP^{X \leftrightarrow Y}$-complexity is at most $\mathsf{poly}\text{-}\log(n)$.

Obviously, $PNP^{X \leftrightarrow Y}_{\square} \subseteq PNP^{X \leftrightarrow Y}$. The question whether the two complexity classes are equal will require our further attention: In particular, the assumption $[|\Gamma_f| \leq 2^{k(n)}]$ doesn't have immediate implications regarding the number of possible accepting sets of rectangles in the (assumed) $P^{X \leftrightarrow Y}$-protocols for f_i-s; what is more, there doesn't have to exist an efficient witness that $[\{ i | f_i(x,y) = \top \} = s]$ as long as $|s|$ is large (say, $n^{\Omega(1)}$) – in contrast to the case of $PNP^{X \leftrightarrow Y}_{\square}$, where the corresponding witness would be the intersection of $|s|$ rectangles, thus itself a rectangle. See Section 3.2 (Lemma 2 in particular).

3.1 Rectangle-patterned non-determinism ($PNP^{X \leftrightarrow Y}_{\square}$) vs. determinism ($P^{X \leftrightarrow Y}$)

Are the complexity classes $PNP^{X \leftrightarrow Y}_{\square}$ and $P^{X \leftrightarrow Y}$ equal?

Lemma 1. *For every total Boolean $f : \mathcal{A} \times \mathcal{B} \to \{\top, \bot\}$,*

$$P^{X \leftrightarrow Y}(f) \in O\left(PNP^{X \leftrightarrow Y}_{\square}(f)^2 \right).$$

Accordingly, $PNP^{X \leftrightarrow Y}_{\square} = P^{X \leftrightarrow Y}$.

Proof. Let Π_f be an $NP^{X \leftrightarrow Y}$-protocol that computes f and witnesses (cf. Definition 4) that $PNP^{X \leftrightarrow Y}_{\square}(f) \leq k(n)$. Let R_Π be the set of Π_f's rectangles and $\Gamma_\Pi \subseteq \mathsf{pow}(R_\Pi)$ be the corresponding family of accepting patterns ($|\Gamma_\Pi| \leq 2^{k(n)}$).

Consider the following $P^{X \leftrightarrow Y}$-protocol Φ for input $(x,y) \in \mathcal{A} \times \mathcal{B}$:

[7] It is not required by the definition, but can be assumed without loss of generality that $m \leq |\Gamma_f|$: the set $\left\{ i_0 \in [m] \middle| \exists\, (x,y) \in \mathcal{A} \times \mathcal{B} : \; f_{i_0}(x,y) = \top, \; \vee_{j \neq i_0} f_j(x,y) = \bot \right\}$ contains at most $|\Gamma_f|$ elements and all other f_i-s are "meaningless" and can be recursively dropped from the decomposition $f(x,y) = \vee_i f_i(x,y)$ without affecting $f(x,y)$ (cf. Footnote 5).

1. $j \leftarrowtail 0$; $\mathscr{A}_1 \leftarrowtail \mathcal{A}$; $\mathscr{B}_1 \leftarrowtail \mathcal{B}$.
2. • $j \leftarrowtail j+1$;
   • $\Gamma_j \leftarrowtail \left\{ \left\{ r \in R_\Pi \middle| (x',y') \in r \right\} \middle| (x',y') \in \mathscr{A}_j \times \mathscr{B}_j \right\}$.
3. If there exists $r_A \times r_B = r \in R_\Pi$ such that
$$\left| \left\{ \gamma \in \Gamma_j \middle| r \in \gamma \right\} \right| \geq \frac{1}{3} \cdot |\Gamma_j| > 0,$$
then do:
   - if $(x,y) \in r$, then output "$\top$" and ***halt***;
   - if $x \notin r_A$, then let $\mathscr{A}_{j+1} \leftarrowtail \mathscr{A}_j \setminus r_A$, else $\mathscr{A}_{j+1} \leftarrowtail \mathscr{A}_j$;
   - if $y \notin r_B$, then let $\mathscr{B}_{j+1} \leftarrowtail \mathscr{B}_j \setminus r_B$, else $\mathscr{B}_{j+1} \leftarrowtail \mathscr{B}_j$;
   - go to Step 2.
4. If there exists $r_A \times r_B = r \in R_\Pi$ such that $x \in r_A$ and
$$\left| \left\{ \gamma \in \Gamma_j \middle| \forall\, x' \in \mathscr{A}_j \cap r_A :\ \exists\, r'_A \times r'_B \in \gamma :\ x' \notin r'_A \right\} \right| \geq \frac{1}{3} \cdot |\Gamma_j| > 0,$$
then do:
   - if $y \in r_B$, then output "$\top$" and ***halt***;
   - $\mathscr{A}_{j+1} \leftarrowtail \mathscr{A}_j \cap r_A$;
   - $\mathscr{B}_{j+1} \leftarrowtail \mathscr{B}_j$;
   - go to Step 2.
5. If there exists $r_A \times r_B = r \in R_\Pi$ such that $y \in r_B$ and
$$\left| \left\{ \gamma \in \Gamma_j \middle| \forall\, y' \in \mathscr{B}_j \cap r_B :\ \exists\, r'_A \times r'_B \in \gamma :\ y' \notin r'_B \right\} \right| \geq \frac{1}{3} \cdot |\Gamma_j| > 0,$$
then do:
   - if $x \in r_A$, then output "$\top$" and ***halt***;
   - $\mathscr{A}_{j+1} \leftarrowtail \mathscr{A}_j$;
   - $\mathscr{B}_{j+1} \leftarrowtail \mathscr{B}_j \cap r_B$;
   - go to Step 2.
6. Output "$\bot$" and ***halt***.

We claim that $\Phi(x,y)$ has complexity $O\big(k(n)^2\big)$ and computes $f(x,y)$.

§1. At the end of Step 2

$$(x,y) \in \mathscr{A}_j \times \mathscr{B}_j$$

always. In the beginning this is trivially true, and the updates (shrinkages) of $\mathscr{A}$ and $\mathscr{B}$ in Steps 3, 4 and 5 occur under conditions that guarantee that (x,y) stays inside $\mathscr{A}_{j+1} \times \mathscr{B}_{j+1}$.

§2. At the end of Step 2

$$\Gamma_j = \left\{ \left\{ r \in R_\Pi \middle| (x',y') \in r \right\} \middle| (x',y') \in \mathscr{A}_j \times \mathscr{B}_j \right\}$$

always.

§3. Answer "$\top$" is always correct. Indeed, producing such an answer necessarily represents having found $r \in R_\Pi$ such that $(x, y) \in r$, therefore $f(x, y) = \top$.

§4. Answer "$\bot$" is always correct. Let j_0 be the value of the index j when "$\bot$" has been produced at Step 6 and assume towards contradiction that $f(x, y) = \top$ and therefore $(x, y) \in r_A \times r_B \in R_\Pi$. If $|\Gamma_{j_0}| = 0$, then the desired contradiction follows readily from §§ 1 and 2, so assume that $|\Gamma_{j_0}| > 0$.

As the entry condition of Step 3 was unsatisfied, it must be the case that

$$\left|\left\{\gamma \in \Gamma_{j_0} \middle| r_A \times r_B \in \gamma\right\}\right| < \frac{1}{3} \cdot \left|\Gamma_{j_0}\right|,$$

and therefore $|\widetilde{\Gamma_{j_0}}| > 2/3 \cdot |\Gamma_{j_0}|$ for

$$\widetilde{\Gamma_{j_0}} \stackrel{\text{def}}{=} \left\{\gamma \in \Gamma_{j_0} \middle| r_A \times r_B \notin \gamma\right\}.$$

Due to §2, $\forall\, \gamma_0 \in \widetilde{\Gamma_{j_0}}$, $\forall\, (x', y') \in \mathscr{A}_{j_0} \times \mathscr{B}_{j_0}$:

$$\left(\forall\, r' \in \gamma_0 :\ (x', y') \in r'\right) \implies (x', y') \notin r_A \times r_B,$$

that is,

$$(x', y') \in r_A \times r_B \implies \exists\, r'_A \times r'_B \in \gamma_0 :\ x' \notin r'_A \vee y' \notin r'_B.$$

As (x', y') can be any pair from the product set $\mathscr{A}_{j_0} \times \mathscr{B}_{j_0}$, the above readily decomposes into

$$\forall\, x' \in \mathscr{A}_{j_0} \cap r_A :\ \exists\, r'_A \times r'_B \in \gamma_0 :\ x' \notin r'_A \bigvee \forall\, y' \in \mathscr{B}_{j_0} \cap r_B :\ \exists\, r'_A \times r'_B \in \gamma_0 :\ y' \notin r'_B.$$

In other words, $\widetilde{\Gamma_{j_0}} = \widetilde{\Gamma^A_{j_0}} \cup \widetilde{\Gamma^B_{j_0}}$ (not necessarily disjointly), where

$$\widetilde{\Gamma^A_{j_0}} \stackrel{\text{def}}{=} \left\{\gamma \in \Gamma_{j_0} \middle| \forall\, x' \in \mathscr{A}_{j_0} \cap r_A :\ \exists\, r'_A \times r'_B \in \gamma :\ x' \notin r'_A\right\}$$

and

$$\widetilde{\Gamma^B_{j_0}} \stackrel{\text{def}}{=} \left\{\gamma \in \Gamma_{j_0} \middle| \forall\, y' \in \mathscr{B}_{j_0} \cap r_B :\ \exists\, r'_A \times r'_B \in \gamma :\ y' \notin r'_B\right\}.$$

As $|\widetilde{\Gamma_{j_0}}| > 2/3 \cdot |\Gamma_{j_0}|$, it necessarily holds that $|\widetilde{\Gamma^A_{j_0}}| > 1/3 \cdot |\Gamma_{j_0}|$ or $|\widetilde{\Gamma^B_{j_0}}| > 1/3 \cdot |\Gamma_{j_0}|$, and therefore the entry condition of Step 4 or 5 must have been satisfied, contradicting our assumption that "$\bot$" was produced at Step 6.

§5. The protocol makes $O(k(n))$ iterations. Due to §1, it is guaranteed by the entry conditions and the actions of Steps 3, 4 and 5 that

$$\left|\Gamma_j\right| \le \left|\Gamma_{j-1}\right| \cdot 2/3$$

at every protocol round $j > 1$. And we have assumed that $|\Gamma_\Pi| \le 2^{k(n)}$.

§6. The $P^{X \leftrightarrow Y}$-complexity of one iteration of the protocol is in $O(k(n))$. As the protocol proceeds, both players locally keep track of j, Γ_j, $\mathscr{A}_j$ and $\mathscr{B}_j$. Non-trivial are only Steps 3, 4 and 5, and it is easy to see that their entry conditions can be checked locally by at least one of the players, and the actions (basically, checking whether $(x, y) \in r = r_A \times r_B$) require $O(\log(|R_\Pi|))$ bits of communication (the cost of sending a "pointer" to $r \in R_\Pi$), which can be assumed to be in $O(k(n))$ (cf. Footnote 5). ■*Lemma 1*

3.2 Patterned non-determinism ($PNP^{X\leftrightarrow Y}$) vs. determinism ($P^{X\leftrightarrow Y}$)

As mentioned earlier, the communication model $PNP^{X\leftrightarrow Y}$ is an interesting object of study because, in particular, the transition from its "rectangular" version $PNP^{X\leftrightarrow Y}_{\square}$ to the general case looks challenging.[8]

Lemma 2. *For every total Boolean $f : \mathcal{A}\times\mathcal{B}\to\{\top,\bot\}$,*

$$P^{X\leftrightarrow Y}(f)\in O\Big(PNP^{X\leftrightarrow Y}(f)^6\Big).$$

Accordingly, $PNP^{X\leftrightarrow Y} = PNP^{X\leftrightarrow Y}_{\square} = P^{X\leftrightarrow Y}$.

To prove it we will use the following simple "hitting set" statement.

Claim 1. *Let $\Gamma\subseteq\mathsf{pow}([m])$ and $t\in\mathbb{N}$ be such that*

$$\forall\,\gamma\in\Gamma:\ |\gamma|\le 2t.$$

Then there exists $\sigma\subseteq[m]$ such that

$$\max_{\gamma\in\Gamma}\big\{|\gamma\cap\sigma|\big\}\le 8e+\log|\Gamma|$$

and

$$\Big|\Big\{\gamma\in\Gamma\Big||\gamma|\ge t,\,\gamma\cap\sigma\neq\emptyset\Big\}\Big|\ge\frac{1}{2}\cdot\Big|\Big\{\gamma\in\Gamma:|\gamma|\ge t\Big\}\Big|.$$

Proof. Let $\sigma_0\in\binom{[m]}{2m/t}$, then

$$\forall\,\gamma\in\Gamma,|\gamma|\ge t:\ \mathop{\boldsymbol{Pr}}_{\sigma_0}\big[\gamma\cap\sigma_0=\emptyset\big]\le\left(\frac{m-t}{m}\right)^{2m/t}=\left(1-\frac{t}{m}\right)^{2m/t}<\frac{1}{4}$$

and

$$\mathop{\boldsymbol{E}}_{\sigma_0}\left[\left|\Big\{\gamma\in\Gamma:|\gamma|\ge t,\,\gamma\cap\sigma_0=\emptyset\Big\}\right|\right]<\frac{1}{4}\cdot\Big|\Big\{\gamma\in\Gamma:|\gamma|\ge t\Big\}\Big|,$$

so

$$\mathop{\boldsymbol{Pr}}_{\sigma_0}\left[\left|\Big\{\gamma\in\Gamma:|\gamma|\ge t,\,\gamma\cap\sigma_0\neq\emptyset\Big\}\right|\ge\frac{1}{2}\cdot\Big|\Big\{\gamma\in\Gamma:|\gamma|\ge t\Big\}\Big|\right]>\frac{1}{2}.\tag{1}$$

Denote $s\stackrel{\text{def}}{=}8e+\log|\Gamma|$. As $\forall\,\gamma\in\Gamma:\ |\gamma|\le 2t$,

$$\forall\,\gamma\in\Gamma:\ \mathop{\boldsymbol{Pr}}_{\sigma_0}\big[|\gamma\cap\sigma_0|\ge s\big]\le\binom{2t}{s}\cdot\left(\frac{2m/t}{m}\right)^s\le\left(\frac{2et}{s}\cdot\frac{2}{t}\right)^s=\left(\frac{4e}{s}\right)^s\le\frac{2^{-8e}}{|\Gamma|},$$

that is

$$\mathop{\boldsymbol{Pr}}_{\sigma_0}\big[\exists\,\gamma\in\Gamma:\ |\gamma\cap\sigma_0|\ge s\big]\le 2^{-8e}.$$

Together with (1) this implies the result. ■*Claim 1*

[8] E.g., if we look at the protocol $\Phi(x,y)$ from Section 3.1, then, first of all, it is not clear how to efficiently generalise for the case of $PNP^{X\leftrightarrow Y}$ the entry conditions of Steps 4 and 5; what is more, the logic underlying those conditions (as represented by the analysis of Φ) doesn't seem to generalise readily.

Proof of Lemma 2. Let

$$f(x,y) \equiv \bigvee_{i=1}^{m} f_i(x,y)$$

be a decomposition that witnesses (cf. Definition 6) that $PNP^{X\leftrightarrow Y}(f) \leq k(n)$. Let $\Gamma_f \subseteq \mathsf{pow}([m])$ be the corresponding family of accepting patterns ($|\Gamma_f| \leq 2^{k(n)}$). Assume without loss of generality that $m \leq |\Gamma_f|$ (cf. Footnote 7). For every $i \in [m]$, let Π_i be a $P^{X\leftrightarrow Y}$-protocol of complexity at most $k(n)$ that computes $f_i(x,y)$ and let R_i be the set of Π_i's *accepting* rectangles ($|R_i| \leq 2^{k(n)}$). Denote for any non-empty $s \subseteq [m]$:

$$R_s \stackrel{\text{def}}{=} \left\{ r_1 \cap \cdots \cap r_{|s|} \middle| r_l \in R_{s(l)} \text{ for } 1 \leq l \leq |s| \right\},$$

that is, R_s is the family of rectangle intersections – therefore rectangles themselves – that witness $[\bigwedge_{i\in s} f_i(x,y)]$ for $(x,y) \in \mathcal{A}\times\mathcal{B}$. Clearly, $|R_s| \leq 2^{|s|\cdot k(n)}$.

Consider the following $P^{X\leftrightarrow Y}$-protocol Ψ for input $(x,y) \in \mathcal{A}\times\mathcal{B}$:

1. $j \leftarrowtail 0$; $\mathscr{A}_1 \leftarrowtail \mathcal{A}$; $\mathscr{B}_1 \leftarrowtail \mathcal{B}$; $s_0 \leftarrowtail \max_{\gamma\in\Gamma_f}\{|\gamma|\}$; $\sigma_0 \leftarrowtail \emptyset$.
2. - $j \leftarrowtail j+1$;
   - $\Gamma_j \leftarrowtail \left\{ \{i | f_i(x',y') = \top\} \middle| (x',y') \in \mathscr{A}_j \times \mathscr{B}_j \right\}$.
3. If $\Gamma_j = \{\emptyset\}$, then output "$\perp$" and ***halt***.
4. If $\{\gamma \cap \sigma_{j-1} | \gamma \in \Gamma_j\} = \{\emptyset\}$, then do:
   - if $\Gamma_j \cap [s_{j-1}/2,\ s_{j-1}] = \emptyset$, then let $s_j \leftarrowtail \max_{\gamma\in\Gamma_j}\{|\gamma|\}$, else $s_j \leftarrowtail s_{j-1}$;
   - let $\sigma_j \subseteq [m]$ be (as guaranteed by Claim 1) such that

     $$\left| \left\{ \gamma \in \Gamma_j \cap [s_j/2,\ s_j] \middle| \gamma \cap \sigma_j \neq \emptyset \right\} \right| \geq \frac{1}{2} \cdot \left| \Gamma_j \cap [s_j/2,\ s_j] \right|$$

     and

     $$\max_{\gamma\in\Gamma_j} \left\{ \left| \gamma \cap \sigma_j \right| \right\} \leq 8e + \log|\Gamma_j|;$$

   else:

   - $s_j \leftarrowtail s_{j-1}$;
   - $\sigma_j \leftarrowtail \sigma_{j-1}$.
5. - $t_j \leftarrowtail \max_{\gamma\in\Gamma_j}\{|\gamma \cap \sigma_j|\}$;
   - $\Delta_j \leftarrowtail \{\gamma \cap \sigma_j : \gamma \in \Gamma_j,\ |\gamma \cap \sigma_j| = t_j\}$;
   - $\Pi_j \leftarrowtail \{\mathscr{A}_j \times \mathscr{B}_j \cap r | r \in R_\delta \text{ for } \delta \in \Delta_j\} \setminus \{\emptyset\}$;
   - $\Pi_j^A \leftarrowtail \left\{ r_A \times r_B \in \Pi_j : \left| \{ r'_A \times r'_B \in \Pi_j : r_A \cap r'_A = \emptyset \} \right| \geq \frac{|\Pi_j|-1}{2} \right\}$;
   - $\Pi_j^B \leftarrowtail \left\{ r_A \times r_B \in \Pi_j : \left| \{ r'_A \times r'_B \in \Pi_j : r_B \cap r'_B = \emptyset \} \right| \geq \frac{|\Pi_j|-1}{2} \right\}$.
6. If there exists $r_A \times r_B \in \Pi_j^A$ such that $x \in r_A$, then do:
   - if $y \in r_B$, then output "$\top$" and ***halt***;
   - $\mathscr{A}_{j+1} \leftarrowtail \mathscr{A}_j \cap r_A$;
   - $\mathscr{B}_{j+1} \leftarrowtail \mathscr{B}_j \setminus r_B$;
   - go to Step 2.
7. If there exists $r_A \times r_B \in \Pi_j^B$ such that $y \in r_B$, then do:

- if $x \in r_A$, then output "$\top$" and ***halt***;
- $\mathscr{A}_{j+1} \leftarrowtail \mathscr{A}_j \setminus r_A$;
- $\mathscr{B}_{j+1} \leftarrowtail \mathscr{B}_j \cap r_B$;
- go to Step 2.

8. - $\mathscr{A}_{j+1} \leftarrowtail \mathscr{A}_j \setminus \bigcup_{r_A \times r_B \in \Pi_j^A} r_A$;
   - $\mathscr{B}_{j+1} \leftarrowtail \mathscr{B}_j \setminus \bigcup_{r_A \times r_B \in \Pi_j^B} r_B$;
   - go to Step 2.

We claim that $\Psi(x, y)$ has complexity $O\big(k(n)^6\big)$ and computes $f(x, y)$.

In the following analysis we call a j-indexed value (e.g., s_j or σ_j) *unchanged* as long as the next round's value is the same as the last round's (e.g., due to assignments like $s_j \leftarrowtail s_{j-1}$ or $\sigma_j \leftarrowtail \sigma_{j-1}$). Otherwise we will say that the corresponding value *changes* at round j.

§1. At the end of Step 2

$$(x, y) \in \mathscr{A}_j \times \mathscr{B}_j$$

always. In the beginning this is trivially true. The updates (shrinkages) of $\mathscr{A}$ and $\mathscr{B}$ in Steps 6 and 7 occur under conditions that guarantee that (x, y) stays inside $\mathscr{A}_{j+1} \times \mathscr{B}_{j+1}$. The updates in Step 8 occur only if the entry conditions of both Steps 6 and 7 were unsatisfied, which also guarantees that (x, y) stays inside.

§2. At the end of Step 2

$$\Gamma_j = \left\{ \left\{ i \middle| f_i(x', y') = \top \right\} \middle| (x', y') \in \mathscr{A}_j \times \mathscr{B}_j \right\}$$

always.

§3. Answer "$\top$" is always correct. Indeed, producing such an answer necessarily represents having found $r \in R_s$ for some non-empty $s \subseteq [m]$ such that $(x, y) \in r$ – that is, the input pair is inside a non-empty intersection of Π_i's accepting rectangles, so $f(x, y) = f_i(x, y) = \top$.

§4. Answer "$\bot$" is always correct. Answering "$\bot$" in Step 3 is conditioned upon $[\Gamma_j = \{\emptyset\}]$, due to §§ 1 and 2 this implies that $f(x, y) = \bot$.

§5. The value of s_j changes at most $\log|\Gamma_f|$ times. This can only happen in Step 4 if the condition $[\Gamma_j \cap [s_{j-1}/2,\ s_{j-1}] = \emptyset]$ is satisfied. We have $s_0 = \max_{\gamma \in \Gamma_f}\{|\gamma|\} \leq m \leq |\Gamma_f|$ and every time s_j changes, it is necessarily the case both that $s_j < s_{j-1}/2$ and $s_j > 0$ (the latter is due to the check in Step 3).

§6. While s_j remains unchanged, σ_j changes at most $\log|\Gamma_f| + 1$ times. The change can happen only in Step 4 if the condition $\left[\left\{\gamma \cap \sigma_{j-1} \middle| \gamma \in \Gamma_j\right\} = \{\emptyset\}\right]$ is satisfied. As long as the value of s_j remains unchanged, every redefinition of σ_j results in $\left\{\gamma \cap \sigma_j \middle| \gamma \in \Gamma_j\right\}$ containing at least half of $\Gamma_j \cap [s_{j-1}/2,\ s_{j-1}]$, and if Γ_j changes, then its content necessarily shrinks – accordingly, there can be at most $\log|\Gamma_j| + 1 \leq \log|\Gamma_f| + 1$ redefinitions of σ_j for the same value of s_j.

§7. While σ_j and s_j remain unchanged, the value of t_j either remains unchanged or decreases (Step 5); $t_j \leq 8e + \log|\Gamma_f|$ always.

§8. While σ_j, s_j and t_j remain unchanged, the protocol makes $O\big(k(n)^2\big)$ iterations. Intuitively, in this situation our protocol solves with respect to $(x, y) \in \mathscr{A}_j \times \mathscr{B}_j$ the problem

$$\text{accept if } \left|\left\{i \in \sigma_j \middle| f_i(x, y) = \top\right\}\right| = t_j,$$

while it is guaranteed by definition that

$$\forall\,(x', y') \in \mathscr{A}_j \times \mathscr{B}_j : \left|\left\{i \in \sigma_j \middle| f_i(x', y') = \top\right\}\right| \leq t_j. \tag{2}$$

Then Δ_j is the set of accepting patterns (in the sense analogous to Definition 5, but with "$\bigvee$" replaced by t_j-threshold) and Π_j is the corresponding family of witnessing rectangle intersections, therefore rectangles themselves. As the updates only can shrink the sets $\mathscr{A}_j$ and $\mathscr{B}_j$, also the family Π_j only shrinks while σ_j, s_j and t_j remain unchanged.

We claim that $\Pi_j = \Pi_j^A \cup \Pi_j^B$ (not necessarily disjointly). Towards contradiction, assume the opposite and let $r_A \times r_B \in \Pi_j \setminus \Pi_j^A \setminus \Pi_j^B$, then there exists $r'_A \times r'_B \in \Pi_j$ such that $r'_A \times r'_B \neq r_A \times r_B$ but $r'_A \times r'_B \cap r_A \times r_B \neq \emptyset$. Let $(x_0, y_0) \in r'_A \times r'_B \cap r_A \times r_B$, $r_A \times r_B = \mathscr{A}_j \times \mathscr{B}_j \cap r$ and $r'_A \times r'_B = \mathscr{A}_j \times \mathscr{B}_j \cap r'$ for $r \in R_\delta$, $r' \in R_{\delta'}$ and δ, $\delta' \in \Delta_j$. If $\delta = \delta'$, then r and r' are distinct elements of

$$R_\delta = \left\{r_1 \cap \cdots \cap r_{|s|} \middle| r_l \in R_{\delta(l)} \text{ for } 1 \leq l \leq |\delta|\right\},$$

contradicting the assumption that each R_i is the set of accepting rectangles in a *deterministic* protocol (whose rectangles are therefore disjoint). If, on the other hand, $\delta \neq \delta'$, then $\delta \cup \delta'$ is a subset of

$$\left\{i \in \sigma_j \middle| f_i(x_0, y_0) = \top\right\},$$

contradicting (2), as $|\delta \cup \delta'| > t_j$. So,

$$\Pi_j = \Pi_j^A \cup \Pi_j^B. \tag{3}$$

Now assume that at round $j+1$ the values of σ_{j+1}, s_{j+1} and t_{j+1} remain unchanged. There are cases to consider.

If the instruction "go to Step 2" has been performed at Step 6 of round j (the case of Step 7 is similar), then $x \in r_A$ such that

$$\left|\left\{r'_A \times r'_B \in \Pi_j : r_A \cap r'_A = \emptyset\right\}\right| \geq \frac{|\Pi_j| - 1}{2}.$$

As $\Pi_{j+1} = \left\{\mathscr{A}_{j+1} \times \mathscr{B}_{j+1} \cap r \middle| r \in R_\delta,\ \delta \in \Delta_{j+1}\right\} \setminus \{\emptyset\}$, the assignment $\mathscr{A}_{j+1} \leftarrowtail \mathscr{A}_j \cap r_A$ at Step 6 of round j has "removed" at least $(|\Pi_j|-1)/2$ elements from Π_{j+1} in comparison to Π_j, and the assignment $\mathscr{B}_{j+1} \leftarrowtail \mathscr{B}_j \setminus r_B$ has removed at least one more (as $r_A \times r_B \in \Pi_j$). Overall, $|\Pi_{j+1}| \leq |\Pi_j|/2$.

If, on the other hand, "go to Step 2" has been performed at Step 8 of round j, then the preceding assignments $\mathscr{A}_{j+1} \leftarrowtail \mathscr{A}_j \setminus \bigcup_{r_A \times r_B \in \Pi_j^A} r_A$ and $\mathscr{B}_{j+1} \leftarrowtail \mathscr{B}_j \setminus \bigcup_{r_A \times r_B \in \Pi_j^B} r_B$ guarantee – along with (3) – that $|\Pi_{j+1}| = \emptyset$ (which, in fact, contradicts our assumption that σ_j, s_j and t_j remain unchanged at round $j+1$).

Accordingly, while σ_j, s_j and t_j remain unchanged since round j_0, the protocol can make only $O(\log|\Pi_{j_0}|)$ iterations. Since $\Delta_{j_0} \subseteq \binom{[m]}{t_{j_0}}$ and $|\Pi_{j_0}| \leq \sum_{\delta \in \Delta_{j_0}} |R_\delta|$ by definition, it holds – as required

– that

$$\left|\Pi_{j_0}\right| \leq 2^{t_{j_0}\cdot k(n)} \cdot \left|\Delta_{j_0}\right| \leq 2^{O\left(k(n)^2\right)},$$

as $t_{j_0} \leq 8e + \log|\Gamma_{j_0}|$ and $|\Delta_{j_0}| \leq |\Gamma_{j_0}| \leq |\Gamma_f| \leq 2^{k(n)}$.

§9. The protocol makes $O\left(k(n)^5\right)$ iterations. Follows readily from §§ 5, 6, 7 and 8.

§10. The $P^{X\leftrightarrow Y}$-complexity of one iteration of the protocol is in $O(k(n))$. To agree upon the value of σ_j in Step 4, the players can, for instance, always pick the lexicographically first suitable candidate (which can be done locally as long as Γ_j and s_j are known to both players). The rest is very similar to the case of $\Phi(x,y)$ in the proof of Lemma 1. ■*Lemma 2*

3.3 Efficient exhaustive witness-searching in $PNP^{X\leftrightarrow Y}$

Assume that certain communication complexity subclass $\mathscr{C} \subseteq NP^{X\leftrightarrow Y}$ is inside $P^{X\leftrightarrow Y}$ for the case of total functions (but not, in general, for the partial-functions case): as discussed earlier, some examples of such subclasses are $NP^{X\leftrightarrow Y} \cap coNP^{X\leftrightarrow Y}$, $UP^{X\leftrightarrow Y}$, $FewP^{X\leftrightarrow Y}$ and $PNP^{X\leftrightarrow Y}$. Let $f : \mathcal{A}\times\mathcal{B} \to \{\top,\bot\}$ belong to $\mathscr{C}$, does it necessarily follow that $NP^{X\leftrightarrow Y}$-witnesses for every $(x,y) \in f^{-1}(\top)$ can be efficiently found?

The answer depends on the precise notion of witness that we have in mind. In particular, as $f \in \mathscr{C} \subseteq P^{X\leftrightarrow Y}$, there is an efficient deterministic protocol that computes f and the transcript of that protocol on any input (x_0,y_0) "witnesses" the value of $f(x_0,y_0)$.

On the other hand, we may consider a specific "canonical" $NP^{X\leftrightarrow Y}$-protocol Π' for f that witnesses the membership $f \in \mathscr{C}$ (recall that $\mathscr{C}$ is a subclass of $NP^{X\leftrightarrow Y}$) – is it necessarily the case that finding a valid Π'-witness for every $(x,y) \in f^{-1}(\top)$ can be done efficiently?

It is so indeed for the cases of $UP^{X\leftrightarrow Y}$ and $FewP^{X\leftrightarrow Y}$: if Π' computes f with at most poly-log(n) distinct witnesses, then a Π'-witness for every $(x,y) \in f^{-1}(\top)$ can be found via essentially the same $P^{X\leftrightarrow Y}$-protocol that is used in the proof of "$FewP^{X\leftrightarrow Y} = P^{X\leftrightarrow Y}$".

It is very similar for $PNP^{X\leftrightarrow Y}$:

Corollary 1. *Let $f : \mathcal{A}\times\mathcal{B} \to \{\top,\bot\}$ be of $PNP^{X\leftrightarrow Y}$-complexity at most $k(n)$, as witnessed via the decomposition $f(x,y) = \vee_{i=1}^{m} f_i(x,y)$.*

Then there exists a deterministic protocol of cost $O\left(k(n)^6\right)$ that receives an input pair $(x_0,y_0) \in \mathcal{A}\times\mathcal{B}$ such that $f(x_0,y_0) = \top$ and outputs some i_0 such that $f_{i_0}(x_0,y_0) = \top$.

Proof. The protocol Ψ from the proof of Lemma 2 finds such i_0 whenever it outputs "$\top$" (cf. §3 of that proof). ■*Corollary 1*

The situation is different in the case of $\mathscr{C} = NP^{X\leftrightarrow Y} \cap coNP^{X\leftrightarrow Y}$ (it was shown by Aho, Ullman and Yannakakis [AUY83] that the class was equal to $P^{X\leftrightarrow Y}$). In [Gav20] subsets $\mathcal{A},\mathcal{B} \subseteq \binom{[n]}{n^{3/5}}$ were presented such that $\forall x \in \mathcal{A}, y \in \mathcal{B} :\ x \cap y \neq \emptyset$, but finding an element from the intersection in a uniformly-random pair $(x,y) \in \mathcal{A}\times\mathcal{B}$ required a randomised communication protocol of complexity $\Omega\left(\sqrt[5]{n}\right)$. If we define $\mathcal{A}' \stackrel{\text{def}}{=} \mathcal{A} \cup \{\emptyset\}$, this will result in a (somewhat) non-trivial instance of the *set intersection problem* with respect to $(x,y) \in \mathcal{A}'\times\mathcal{B}$. We can apply the following reasoning:

- the problem is in $NP^{X\leftrightarrow Y}$: denote by Π the $NP^{X\leftrightarrow Y}$-protocol that accepts (x,y) if and only if it receives some $i \in x \cap y$ as a witness;

- the problem is in $coNP^{X\leftrightarrow Y}$, as $x \cap y = \emptyset$ only happens when $x = \emptyset$ and this condition can be easily checked even without a witness;
- due to [AUY83], the corresponding set intersection problem is in $NP^{X\leftrightarrow Y} \cap coNP^{X\leftrightarrow Y} = P^{X\leftrightarrow Y}$;
- nevertheless, given an input pair (x, y) such that $x \cap y \neq \emptyset$, finding a valid witness for $\Pi(x, y)$ cannot be done efficiently (even by a randomised protocol).

Can we strengthen Lemma 2 and Corollary 1 even further? Namely, if the membership $f \in PNP^{X\leftrightarrow Y}$ is established via considering the decomposition $f(x, y) = \vee_{i=1}^{m} f_i(x, y)$, can we efficiently find for every given input $(x, y) \in f^{-1}(\top)$ the *exhaustive* list of the corresponding $NP^{X\leftrightarrow Y}$-witnesses, that is, the exact content of $\left\{i \middle| f_i(x, y) = \top\right\}$?

Note that a positive answer wouldn't follow trivially from the repeated application of the efficient witness-finding protocol of Corollary 1: unlike $FewP^{X\leftrightarrow Y}$, $PNP^{X\leftrightarrow Y}$ allows arbitrarily large sets of witnesses for the same $(x, y) \in f^{-1}(\top)$, and therefore such repeated application until all valid witnesses are exhausted can be inefficient. On the other hand, efficiency considerations do not readily lead to the negative answer either: Π is a $PNP^{X\leftrightarrow Y}$-protocol and therefore it admits at most $2^{\text{poly-log}(n)}$ different *patters* (that is, possible exhaustive sets of valid witnesses) – accordingly, a deterministic protocol of complexity poly-log(n) can have enough distinct "leaves" for returning every answer at least once (each protocol leaf is marked by the corresponding answer and every possible pattern must be the answer corresponding to at least one leaf).

As f itself is less relevant for the problem of *exhaustive* witness-searching, we are switching to the functions f_i as our primary objects of concern.

Theorem 1 (*Efficient exhaustive $PNP^{X\leftrightarrow Y}$-search*)**.** *Let $f_1(x, y), \ldots, f_m(x, y) : \mathcal{A} \times \mathcal{B} \to \{\top, \bot\}$ be such that $\forall\, i :\ P^{X\leftrightarrow Y}(f_i) \leq k(n)$ and*

$$\left|\left\{\left\{i \in [m] \middle| f_i(x, y) = \top\right\} : (x, y) \in \mathcal{A} \times \mathcal{B}\right\}\right| \leq 2^{\ell(n)}.$$ [9]

Then there exists a deterministic protocol of cost $O\big(k(n)^6 \cdot \ell(n)\big)$ that receives an input pair $(x_0, y_0) \in \mathcal{A} \times \mathcal{B}$ and outputs the set $\left\{i \middle| f_i(x_0, y_0) = \top\right\}$.

That is, the *exhaustive-search function* $F(x, y) \stackrel{\text{def}}{=} \left\{i \middle| f_i(x, y) = \top\right\}$ is in $P^{X\leftrightarrow Y}$ if $f(x, y) \equiv \vee_{i=1}^{m} f_i(x, y)$ is a bipartite total function in $PNP^{X\leftrightarrow Y}$. The statement is optimal from the structural perspective: if f_i-s were not in $P^{X\leftrightarrow Y}$, then F would not be there either; on the other hand, any deterministic protocol for F must have at least $\left|\left\{\left\{i \middle| f_i(x, y) = \top\right\}\right\}_{x,y}\right|$ leaves (as discussed above).

Proof. Consider the following $P^{X\leftrightarrow Y}$-protocol Ξ for input $(x, y) \in \mathcal{A} \times \mathcal{B}$:

1. $j \leftarrowtail 0$; $\Gamma_1 \leftarrowtail \left\{\left\{i \in [m] \middle| f_i(x, y) = \top\right\} : (x, y) \in \mathcal{A} \times \mathcal{B}\right\}$.
2. • $j \leftarrowtail j + 1$;
   • $\mathscr{W}_j^{+} \leftarrowtail \left\{i \in [m] \middle| \left|\left\{\gamma \in \Gamma_j \middle| i \in \gamma\right\}\right| \geq |\Gamma_j|/2\right\}$;
   • $\mathscr{W}_j^{-} \leftarrowtail [m] \setminus \mathscr{W}_j^{+}$.
3. If for some $i_0 \in \mathscr{W}_j^{-}$ it holds that $f_{i_0}(x, y) = \top$, then do:
   • $\Gamma_{j+1} \leftarrowtail \left\{\gamma \in \Gamma_j \middle| i_0 \in \gamma\right\}$;
   • go to Step 2.

[9] Here m can be any function of n, cf. Footnotes 5 and 7.

4. If for some $i_0 \in \mathscr{W}_j^+$ it holds that $f_{i_0}(x,y) = \bot$, then do:
   - $\Gamma_{j+1} \leftarrowtail \{\gamma \in \Gamma_j \,|\, i_0 \notin \gamma\}$;
   - go to Step 2.
5. Output W_j^+ and ***halt***.

We claim that $\Xi(x,y)$ has complexity $O\big(k(n)^6 \cdot \ell(n)\big)$ and outputs the set $F(x,y) = \big\{i \,\big|\, f_i(x,y) = \top\big\}$.

§1. At the end of Step 2

$$\left\{i \,\middle|\, f_i(x,y) = \top\right\} \in \Gamma_j$$

always. In the beginning this is trivially true, and the updates (shrinkages) of Γ in Steps 3 and 4 occur under conditions that guarantee that $\big\{i \,\big|\, f_i(x,y) = \top\big\}$ stays inside Γ_{j+1}.

§2. The answer is always correct. If the conditions of steps both 3 and 4 where unsatisfied, then it must be the case that

$$W_j^+ = \left\{i \,\middle|\, f_i(x,y) = \top\right\}$$

in Step 5.

§3. The protocol makes $O(\ell(n))$ iterations. It follows from the definitions of W_j^+ and W_j^- that

$$\big|\Gamma_{j+1}\big| \le |\Gamma_j|/2$$

if the condition of either Step 3 or Step 4 is satisfied, §1 guarantees that $\Gamma_j \neq \emptyset$ and it is assumed by the theorem statement that $|\Gamma_1| \le 2^{\ell(n)}$.

§4. The $P^{X \leftrightarrow Y}$-complexity of one iteration of the protocol is in $O\big(k(n)^6\big)$. To perform the check of Step 3 we use the protocol guaranteed by Corollary 1 with respect to

$$f'(x,y) \stackrel{\text{def}}{=} \vee^m_{i \in W_j^-} f_i(x,y),$$

and for the check in Step 4 we use it with

$$f''(x,y) \stackrel{\text{def}}{=} \vee^m_{i \in W_j^+} \neg f_i(x,y),$$

where $\neg f_i(\cdot,\cdot)$ stands for the negation of the predicate $f_i(\cdot,\cdot)$. Clearly, the functions $f', f'' : \mathcal{A} \times \mathcal{B} \to \{\top, \bot\}$ satisfy the requirements of Corollary 1 and the corresponding i_0 meets the needs of Steps 3 and 4, respectively. ■*Theorem 1*

4 Three-party communication with listening Charlie

Let $g(x,y,z) : \mathcal{A} \times \mathcal{B} \times \mathcal{C} \to \{0,1\}$ be a tripartite total function.[10] The three input values will always be partitioned according to the *"number in hand"* input partition, that is, Alice receives x, Bob receives

[10] We are now using $\{0,1\}$ as the default range in the definitions of the communication problems (as opposed to the previously used $\{\top,\bot\}$) as they no longer possess any "logical asymmetry": in the constructions of Section 3 that resulted from the asymmetry in the standard notion of computational non-determinism. Moreover, the results of this part would remain valid if the considered problem were a *tripartite total function with any range* (the proof of Theorem 2 would be based on the same idea but phrased somewhat differently if $\big|\{g(x,y,z)\}_{x,y,z}\big| \in \omega(1)$).

y and Charlie receives z.

The only three-party communication regime that we will consider in this work is *deterministic*, therefore we will drop "$P^{X \leftrightarrow Y}$" and only depict the "communication layout" of each model in the corresponding notation. For that we will deliberately use a not-too-abbreviated, but hopefully, rather intuitive layout representation (cf. Footnote 4).

Definition 7 ($[(X \leftrightarrow Y) \to Z]$, *interacting Alice and Bob with listening Charlie*)**.** *For $n \in \mathbb{N}$, let the sets $|\mathcal{A}|, |\mathcal{B}|, |\mathcal{C}|$ be such that $\max\{|\mathcal{A}|, |\mathcal{B}|, |\mathcal{C}|\} \in (2^{n-1}, 2^n]$ and let $g : \mathcal{A} \times \mathcal{B} \times \mathcal{C} \to \{0,1\}$.*

Let Π be a deterministic protocol where

- *Alice receives x, Bob receives y and Charlie receives z;*
- *Alice and Bob interact in the broadcasting regime, that is, Charlie receives the transcript of their communication;*
- *Charlie produces the answer.*

If the transcript of $\Pi(x,y,z)$ contains at most $k(n)$ bits and the protocol computes $g(x,y,z)$, then we say that the $[(X \leftrightarrow Y) \to Z]$-complexity of g is at most $k(n)$. We say that g is efficiently computable in $[(X \leftrightarrow Y) \to Z]$ if its complexity in the model is at most $\text{poly-log}(n)$.

Alternatively, the model $[(X \leftrightarrow Y) \to Z]$ can be described as letting Alice and Bob interact in order to produce a message that is sent to Charlie, who must answer upon receiving it.

The following models could be pictured as "merging" a pair of players (or allowing free communication between them). In those cases we will address the merged player by the name of the first individual (e.g., "Alice+Bob" – the player who receives input (x,y) – will be called Alice and so on).

Definition 8 ($[(X,Y) \to Z]$)**.** *An $[(X,Y) \to Z]$-protocol is a deterministic protocol where*

- *Alice receive (x,y); Charlie receives z;*
- *Alice send a message to Charlie;*
- *Charlie produces the answer.*

If it computes $g : \mathcal{A} \times \mathcal{B} \times \mathcal{C} \to \{0,1\}$, then we say that the $[(X,Y) \to Z]$-complexity of g is at most the maximum number of bits sent by this protocol for the given input length. The notion of efficiency for $[(X,Y) \to Z]$ is similar to that for $[(X \leftrightarrow Y) \to Z]$ (cf. Definition 7).

Definition 9 ($[(X,Z) \leftrightarrow (Y,Z)]$)**.** *An $[(X,Z) \leftrightarrow (Y,Z)]$-protocol is a deterministic protocol where*

- *Alice receives (x,z) and Bob receives (y,z);*
- *they interact;*
- *Bob produces the answer.*

If it computes $g : \mathcal{A} \times \mathcal{B} \times \mathcal{C} \to \{0,1\}$, then we say that the $[(X,Z) \leftrightarrow (Y,Z)]$-complexity of g is at most the maximum number of bits sent by this protocol for the given input length. The notion of efficiency for $[(X,Z) \leftrightarrow (Y,Z)]$ is similar to that for $[(X \leftrightarrow Y) \to Z]$ (cf. Definition 7).

Definition 10 ($[(X,Z) \leftrightarrow (Y,Z)] \cap [(X,Y) \to Z]$)**.** *For $n \in \mathbb{N}$, let the sets $|\mathcal{A}|, |\mathcal{B}|, |\mathcal{C}|$ be such that $\max\{|\mathcal{A}|, |\mathcal{B}|, |\mathcal{C}|\} \in (2^{n-1}, 2^n]$ and let $g : \mathcal{A} \times \mathcal{B} \times \mathcal{C} \to \{0,1\}$.*

The $[(X,Z) \leftrightarrow (Y,Z)] \cap [(X,Y) \to Z]$-complexity of g is the maximum of its $[(X,Z) \leftrightarrow (Y,Z)]$- and $[(X,Y) \to Z]$-complexities; if that is at most $\text{poly-log}(n)$, then we say that g is efficiently computable in $[(X,Z) \leftrightarrow (Y,Z)] \cap [(X,Y) \to Z]$.

Recall that the models both $[(X,Z) \leftrightarrow (Y,Z)]$ and $[(X,Y) \to Z]$ can be viewed as *amplifications of* $[(X \leftrightarrow Y) \to Z]$ obtained via letting a pair of players communicate for free; accordingly, solving a

communication problem in $[(X \leftrightarrow Y) \to Z]$ is at least as hard as solving it in $[(X, Z) \leftrightarrow (Y, Z)] \cap [(X, Y) \to Z]$.

Theorem 2. *Let $g : \mathcal{A} \times \mathcal{B} \times \mathcal{C} \to \{0,1\}$ be a tripartite total function whose $[(X, Z) \leftrightarrow (Y, Z)]$-complexity is $k(n)$ and $[(X, Y) \to Z]$-complexity is $\ell(n)$, then the $[(X \leftrightarrow Y) \to Z]$-complexity of g is in $O\big(k(n)^6 \cdot \ell(n)\big)$. In particular, $g(x, y, z)$ has an efficient $[(X \leftrightarrow Y) \to Z]$-protocol if and only if it is efficiently computable in $[(X, Z) \leftrightarrow (Y, Z)] \cap [(X, Y) \to Z]$.*

Proof. As both $[(X, Z) \leftrightarrow (Y, Z)]$ and $[(X, Y) \to Z]$ are amplifications of $[(X \leftrightarrow Y) \to Z]$, the existence of an efficient protocol in the latter model trivially implies efficient computability in $[(X, Z) \leftrightarrow (Y, Z)] \cap [(X, Y) \to Z]$.

Consider an $[(X, Y) \to Z]$-protocol Π_1 and let $\alpha_{x,y} \in \{0,1\}^{\ell(n)}$ denote the message sent by Alice to Charlie when her input is $(x, y) \in \mathcal{A} \times \mathcal{B}$.[11] As receiving this message allows Charlie to compute $g(x, y, z)$, every possible message $\alpha_{x,y}$ corresponds to a function $\mathcal{C} \to \{0,1\}$: this function is the description of Charlie's behaviour when he receives the corresponding message from Alice and $z \in \mathcal{C}$ as input.

For every $z_0 \in \mathcal{C}$, let

$$f_{z_0}(x, y) \stackrel{\text{def}}{=} \begin{cases} \top & \text{if } g(x, y, z_0) = 1; \\ \bot & \text{otherwise.} \end{cases}$$

We will apply Theorem 1 to the family $(f_z)_{z \in \mathcal{C}}$. On the one hand, the $P^{X \leftrightarrow Y}$-complexity of $f_{z_0}(x, y)$ is the $[(X, Z) \leftrightarrow (Y, Z)]$-complexity of computing $g(x, y, z_0)$ when $(x, y) \in \mathcal{A} \times \mathcal{B}$, which is at most $k(n)$. On the other hand, for every $(x_0, y_0) \in \mathcal{A} \times \mathcal{B}$ it holds that

$$\Big\{ z \in \mathcal{C} \Big| f_z(x_0, y_0) = \top \Big\} = \Big\{ z \Big| g(x_0, y_0, z) = 1 \Big\} = \Big\{ z \Big| \alpha_{x_0,y_0}(z) = 1 \Big\},$$

where we let "$\alpha_{x_0,y_0}(\cdot)$" stand for the function in $\mathcal{C} \to \{0,1\}$ that the corresponding message represents, as discussed above. Accordingly, every such set corresponds to some $\alpha_{x,y} \in \{0,1\}^{\ell(n)}$ and

$$\left| \left\{ \left\{ z \in \mathcal{C} \Big| f_z(x_0, y_0) = \top \right\} : (x_0, y_0) \in \mathcal{A} \times \mathcal{B} \right\} \right| \leq 2^{\ell(n)}.$$

Theorem 1 guarantees the existence of a deterministic bipartite protocol of cost $O\big(k(n)^6 \cdot \ell(n)\big)$ that receives $(x, y) \in \mathcal{A} \times \mathcal{B}$ and computes the set $\big\{ z_0 \in \mathcal{C} \big| f_{z_0}(x, y) = \top \big\}$. In our $[(X \leftrightarrow Y) \to Z]$-protocol Alice and Bob will use that procedure, then send to Charlie some $\alpha_{x',y'} \in \{0,1\}^{\ell(n)}$ that corresponds to that set, that is,

$$\Big\{ z_0 \in \mathcal{C} \Big| f_{z_0}(x, y) = \top \Big\} = \Big\{ z_0 \Big| \alpha_{x',y'}(z_0) = 1 \Big\}.$$

Upon receiving it, Charlie, who knows z, will answer with $\alpha_{x',y'}(z) = g(x, y, z)$. ■*Theorem 2*

5 Conclusions

The study of communication complexity was initiated by Abelson [Abe78], it was aimed at "*assessing the complexity of computations carried out in distributed networks*". Since then our understanding of the area has somewhat advanced, so it may be desirable to summarise the achievement and to identify

[11] Recall that we are addressing "merged" players by the name of the first included individual.

new interesting directions. This work has been motivated both by the original question due to P. Hrubeš and by the latter goal.

We saw a new structural result in the context of bipartite communication complexity of total functions: on the one hand, it could be viewed as a rather natural generalisation of what was known previously; on the other hand, it had somewhat non-trivial implications for the multi-party case.[12] The main theme of this work was looking for limitations that were imposed by the assumed *total structure* of the communication problem upon the "structural diversity" of communication complexity classes:

- in the two-party case we've seen that the newly defined class $PNP^{X\leftrightarrow Y}$ – a generalisation of previously studied $UP^{X\leftrightarrow Y}$ and $FewP^{X\leftrightarrow Y}$ – admits efficient deterministic protocols, i.e., $PNP^{X\leftrightarrow Y} \subseteq P^{X\leftrightarrow Y}$;
- in the multi-party case we've seen that the class of tripartite total functions efficiently computable by deterministic interacting Alice and Bob with listening Charlie (the model that we denoted by $[(X \leftrightarrow Y) \rightarrow Z]$) equals the weakest among its two-party amplifications obtained by allowing free communication between a pair of players.

Previously known examples of such limitations are $NP^{X\leftrightarrow Y} \cap coNP^{X\leftrightarrow Y} \subseteq P^{X\leftrightarrow Y}$ [AUY83], $UP^{X\leftrightarrow Y} \subseteq P^{X\leftrightarrow Y}$ [Yan91] and $FewP^{X\leftrightarrow Y} \subseteq P^{X\leftrightarrow Y}$ [KNSW94]. None of these five inclusions (in fact, equalities) among the classes would hold if the functions to be computed were not total.

From the combinatorial standpoint, the case of total functions is very natural in the context of communication complexity. What other structural implications does it have? In particular, what are the "strengths of determinism" that are exclusive for total functions?

Acknowledgements

I am grateful to Pavel Hrubeš both for the original motivating question and for numerous insightful discussions. A number of very useful suggestions, both technical and editorial, have been received from colleagues and anonymous reviewers.

References

[Abe78] H. Abelson. Lower Bounds on Information Transfer in Distributed Computations. *Proceedings of the 19th Annual Symposium on Foundations of Computer Science*, pages 151–158, 1978.

[AUY83] A. Aho, J. Ullman, and M. Yannakakis. On Notions of Information Transfer in VLSI Circuits. *Proceedings of the 15th Symposium on Theory of Computing*, pages 133–139, 1983.

[DKW11] J. Draisma, E. Kushilevitz, and E. Weinreb. Partition arguments in multiparty communication complexity. *Theoretical Computer Science 412(24)*, pages 2611–2622, 2011.

[Gav20] D. Gavinsky. The Communication Complexity of the Inevitable Intersection Problem. *Chicago Journal of Theoretical Computer Science*, article 3, 2020.

[GPW18] M. Göös, T. Pitassi, and T. Watson. The Landscape of Communication Complexity Classes. *Computational Complexity 27(2)*, pages 245–304, 2018.

[12] The three-party construction from Section 4 can possibly be generalised for more participants.

[KN97] E. Kushilevitz and N. Nisan. Communication Complexity. *Cambridge University Press*, 1997.

[KNSW94] M. Karchmer, I. Newman, M. Saks, and A. Wigderson. Non-deterministic Communication Complexity with Few Witnesses. *Journal of Computer and System Sciences 49(2)*, pages 247–257, 1994.

[Yan91] M. Yannakakis. Expressing Combinatorial Optimization Problems by Linear Programs. *Journal of Computer and System Sciences 43(3)*, pages 441–466, 1991.

[Yao79] A. C-C. Yao. Some Complexity Questions Related to Distributive Computing. *Proceedings of the 11th Symposium on Theory of Computing*, pages 209–213, 1979.

ECCC ISSN 1433-8092

https://eccc.weizmann.ac.il