THEORY OF COMPUTING, Volume 17 (8), 2021, pp. 1–28
www.theoryofcomputing.org

The Layer Complexity of Arthur-Merlin-like Communication

Dmitry Gavinsky*

Received May 23, 2019; Revised December 11, 2020; Published September 27, 2021

Abstract. In communication complexity the *Arthur-Merlin (AM)* model is the most natural one that allows both randomness and nondeterminism. Presently we do not have any super-logarithmic lower bound for the *AM*-complexity of an explicit function. Obtaining such a bound is a fundamental challenge to our understanding of communication phenomena.

In this article we explore the gap between the known techniques and the complexity class *AM*. In the first part we define a new natural class, *Small-advantage Layered Arthur-Merlin (SLAM)*, that has the following properties:

- *SLAM* is (strictly) included in *AM* and includes all previously known subclasses of *AM* with non-trivial lower bounds:

$$NP, MA, SBP, UAM \subseteq SLAM \subset AM.$$

  Note that $NP \subset MA \subset SBP$, while *SBP* and *UAM* are known to be incomparable.
- *SLAM* is qualitatively stronger than the union of those classes:

$$f \in SLAM \setminus (SBP \cup UAM)$$

  holds for an (explicit) partial function f.

*Partially funded by the grant 19-27871X of GA ČR. Part of this work was done while visiting the Centre for Quantum Technologies at the National University of Singapore, and was partially supported by the Singapore National Research Foundation, the Prime Minister's Office and the Ministry of Education under the Research Centres of Excellence programme under grant R 710-000-012-135.

ACM Classification: F.1.3

AMS Classification: 68Q15

Key words and phrases: communication complexity, complexity classes

DOI: 10.4086/toc.2021.v017a008

- *SLAM* is a subject to the discrepancy bound: for any f

$$SLAM(f) \in \Omega\left(\sqrt{\log \frac{1}{disc(f)}}\right).$$

In particular, the inner product function does not have an efficient *SLAM*-protocol.

Structurally this can be summarised as

$$SBP \cup UAM \subset SLAM \subseteq AM \cap PP.$$

In the second part we ask why proving a lower bound of $\omega(\sqrt{n})$ on the *MA*-complexity of an explicit function seems to be difficult. We show that such a bound either must explore certain "uniformity" of *MA* (which would require a rather unusual argument), or would imply a non-trivial lower bound on the *AM*-complexity of the same function.

Both of these results are related to the notion of *layer complexity*, which is, informally, the number of "layers of nondeterminism" used by a protocol.

1 Introduction

The communication model *Arthur-Merlin (AM)* is beautiful. It is the most natural regime that allows both randomness and nondeterminism. Informally,

- *BPP*, the canonical complexity class representing *randomised communication*, contains such bipartite functions f that admit an *approximate partition* of the set $f^{-1}(1)$ into quasi-polynomially many rectangles;

- *NP*, the canonical complexity class representing *nondeterministic communication*, contains such f that admit an *exact cover* of the set $f^{-1}(1)$ by quasi-polynomially many rectangles;

- *AM* contains such f that admit an *approximate cover* of the set $f^{-1}(1)$ by quasi-polynomially many rectangles.

While both *BPP* and *NP* are relatively well understood and many strong and tight lower bounds are known, we do not have any non-trivial lower bound for the *AM*-complexity of an explicit function. Obtaining such a bound is a fundamental challenge to our understanding of communication complexity.

Among numerous analytical efforts that have been made to understand *AM*, in this paper we are paying special attention to these two:

- In 2003 Klauck [13] studied the class *Merlin-Arthur (MA)*: while (again, informally) *AM* can be viewed as "randomness over nondeterminism," *MA* is "nondeterminism over randomness." Klauck has found an elegant way to exploit this difference in order to prove strong lower bounds against *MA*.

- In 2015 Göös, Pitassi and Watson [10] demonstrated strong lower bounds against the class *Unambiguous Arthur-Merlin (UAM)*, which was defined in the same paper. Similarly to *AM* (and unlike *MA*), their class can be viewed as "randomness over nondeterminism," but only a very special form of nondeterminism is allowed: namely, only the (erroneously accepted) elements of $f^{-1}(0)$ may belong to several rectangles; every element of $f^{-1}(1)$ can belong to at most one rectangle of the nondeterministic cover. In other words, a *UAM*-protocol must correspond to an approximate partition of $f^{-1}(1)$, but at the same time it may be an arbitrary cover of a small fraction of $f^{-1}(0)$. Intuitively, a *UAM*-protocol must "behave like *BPP*" over $f^{-1}(1)$ and is unrestricted over the small erroneously accepted fraction of $f^{-1}(0)$.

Interestingly, the classes *MA* and *UAM* are *incomparable*: from the lower bounds demonstrated in [10] and in [9] it follows that

$$UAM \not\subseteq MA \quad \text{and} \quad MA \not\subseteq UAM\,.$$

In the first half of this article (Section 3) we try to find a communication model that would be as close to *AM* as possible, while staying within the reach of our analytic abilities. Inspired by the (somewhat Hegelian) metamorphosis of "easy" *BPP* and *NP* into "hard" *AM*, we will try to apply a similar "fusion" procedure to the classes *MA* and *UAM*, hoping that the outcome will give us some new insight into the mystery of *AM*.

Namely, we start by looking for a communication complexity class, defined as naturally as possible and containing both *MA* and *UAM*. We will call it *Layered Arthur-Merlin (LAM)* (Definition 3.3). Informally, it can be described as letting a protocol behave like *MA* over $f^{-1}(1)$ and arbitrarily over the erroneously accepted small fraction of $f^{-1}(0)$. Note that it follows trivially from the previous discussion (at least on the intuitive level) that

$$MA \cup UAM \subseteq LAM\,.$$

Then we will add a few rather technical "enhancements" to *LAM* in order to get a class that includes all previously known classes "under *AM*" with non-trivial lower bounds: most noticeably, the class *SBP*, which is known to be strictly stronger than *MA* and strictly weaker than *AM* (see [11, 9, 14]).

We call the resulting model *Small-advantage Layered Arthur-Merlin (SLAM)* (Definition 3.4) and it holds that

$$MA, UAM, SBP, LAM \;\subseteq\; SLAM \;\subset\; AM\,.$$

Moreover, we will demonstrate a partial function

$$f \in SLAM \setminus (UAM \cup SBP)\,,$$

that is, *SLAM* will be strictly stronger than the union of all subclasses of *AM* with previously known non-trivial lower bounds (as $UAM \cup SBP$ includes them all).[1]

[1] Here and later when referring to the *subclasses of AM with previously known non-trivial lower bounds*,we mean, in particular, the classes that are *known* to be included by *AM*. Note that not only do we not have any non-trivial lower bound against *AM* yet, but we also cannot guarantee that any interesting complexity class is not included in *AM*.

Both *LAM* and *SLAM* seem to require a new approach for proving lower bounds. It will be developed in Section 3.1, showing that these classes are still a subject to the discrepancy bound: for any function f,

$$SLAM(f) \in \Omega\left(\sqrt{\log\frac{1}{\mathrm{disc}(f)}}\right),$$

where $SLAM(f)$ denotes the "*SLAM*-complexity" of f. In particular, the *inner product function* does not have an efficient *SLAM*-protocol.

These properties of *SLAM* can be summarised structurally as

$$SBP \cup UAM \subset SLAM \subseteq AM \cap PP,$$

where *PP* is the class consisting of functions with high discrepancy.

The problem of proving a lower bound of $\omega(\sqrt{n})$ for the *MA*-complexity of an explicit function has been open since 2003, when Klauck [13] showed that the *MA*-complexity of *Disj* and *IP* was in $\Omega(\sqrt{n})$. At that point a number of researchers believed that the actual *MA*-complexity of these problems was in $\Omega(n)$, so it was surprising when Aaronson and Wigderson [1] demonstrated *MA*-protocols for *Disj* and *IP* of cost $O(\sqrt{n}\log n)$, which was later improved by Chen [6] to $O\left(\sqrt{n\log n\log\log n}\right)$.

In the second part of this article (Section 4) we try to understand why proving a super-$\sqrt{n}$ lower bound against *MA* seems to be difficult. We will define a communication model $\widetilde{MA}$ that can be viewed, in certain sense, as a *non-uniform MA*. On the one hand, we will see that imposing the corresponding uniformity constraint on $\widetilde{MA}$-protocols makes them not stronger than *MA*-protocols; on the other hand, all known lower bounds on $MA(f)$ readily translate to similar lower bounds on $\widetilde{MA}(f)$.

Intuitively, a complexity analysis that would explore the uniformity of *MA* (as opposed to $\widetilde{MA}$) must have a very unusual structure: the difference between the classes is subtle and we are not aware of any examples where this type of an argument is used. At the same time, we will see that $\widetilde{MA}(f) \in O\left(\sqrt{n\cdot AM(f)}\right)$ for any function f—that is, any lower bound of the form $\widetilde{MA}(f) \in \omega(\sqrt{n})$ would have non-trivial implications for $AM(f)$. This partially explains why no super-$\sqrt{n}$ lower bound against *MA* has been found yet.[2]

Why *LAM* is interesting. In the hope that it would benefit the reader, let us explain the motivation for defining and studying the communication models presented in the first part of this article. The strong lower bounds that were shown earlier for both *MA* and *UAM* were in the first place *steps towards AM*. Both *MA* and *UAM* have very natural definitions, they both can be viewed as weakened versions of *AM*, and the authors of both [13] and [10] have invented new insightful approaches while analysing these models.

[2]It is relatively easy to show $AM(f) \in \Omega(\log n)$ for an explicit function (see Footnote 10); to improve that, a lower bound of the form $\widetilde{MA}(f) \in \omega\left(\sqrt{n\cdot\log n}\right)$ would be needed. However, it seems that proving any $\widetilde{MA}(f_0) \in \omega(\sqrt{n})$ and deriving from it, via the argument of Section 4, that $AM(f_0) \in \omega(1)$ would by itself shed some new light on the enigma of *AM*. As one of the concluding open problems (Section 5), we suggest proving a lower bound of the form $\Omega\left(\sqrt{n\log n}\right)$ on the *MA*-complexity of an explicit function.

The model *LAM*, in turn, has been defined as a natural "junction" of *MA* and *UAM*, at least as strong as either of the predecessors.[3] As the known approaches for analysing *MA* and *UAM* were rather different qualitatively, we expected the new model to be challenging enough to justify defining it. Our experience of proving a strong lower bound for the newly defined model has confirmed those expectations.

We hope that studying *LAM* will serve as the next step towards understanding *AM*.

2 Preliminaries and definitions

For $x \in \{0,1\}^n$ and $i \in [n] = \{1,\ldots,n\}$, we will write x_i or $x(i)$ to address the i-th bit of x (preferring "x_i" unless it may cause ambiguity). Similarly, for $S \subseteq [n]$, let both x_S and $x(S)$ denote the $|S|$-bit string, consisting of (naturally ordered) bits of x, whose indices are in S.

For a (discrete) set A and $k \in \mathbb{N}$, we denote by $\mathrm{Pow}(A)$ the set of subsets of A and by $\binom{A}{k}$ the set $\{a \in \mathrm{Pow}(A) \mid |a| = k\}$.

Our primary objects of computation will be *bipartite Boolean functions* of the form $A \times B \to \{0,1\}$ (typically, $\{0,1\}^n \times \{0,1\}^n \to \{0,1\}$). At times we will consider *partial bipartite Boolean functions*, where some of the pairs are *excluded*: this can be interpreted either as assuming that those pairs are never given as input, or as allowing any output of a communication protocol when those pairs are received. We will view partial Boolean functions as total ones that are taking values from $\{0,1,\perp\}$, where "$\perp$" marks the excluded input values. Note that the total functions are a special case, so $f : A \times B \to \{0,1,\perp\}$ can be either total or partial. When we refer to *an input distribution for a function* $f : A \times B \to \{0,1,\perp\}$, we mean a distribution defined on $f^{-1}(0) \cup f^{-1}(1)$.

We will use the *logical OR* ($\bigvee$) operator with respect to partial Boolean functions, defined as follows (note the asymmetry between the first two cases):

$$f_1(x) \vee f_2(y) \stackrel{\text{def}}{=} \begin{cases} 1 & \text{if } f_1(x) = 1 \text{ or } f_2(y) = 1\,; \\ 0 & \text{if } f_1(x) = 0 \text{ and } f_2(y) = 0\,; \\ \perp & \text{otherwise.} \end{cases} \tag{2.1}$$

2.1 Communication complexity

We refer to [15] for a classical background on communication complexity and to [11] for a great survey of the more recent developments.

Communication problems. We will repeatedly consider the following two communication problems.

Definition 2.1 (*Disjointness function, Disj*)**.** For every $n \in \mathbb{N}$, let $(x,y) \in \{0,1\}^n \times \{0,1\}^n$. Then

$$Disj(x,y) = \bigwedge_{i=1}^{n} (x_i = 0 \vee y_i = 0)\,.$$

[3] Here we are referring to *LAM*, as its definition is more natural and less technically involved than that of *SLAM*; on the other hand, the difference between the two models is, in our opinion, merely formal (as explained above, we wanted the corresponding complexity class to contain all previously studied subclasses of *AM*, including *SBP*, and that was the reason for "boosting" *LAM*, which resulted in *SLAM*).

Definition 2.2 (*Inner product function, IP*)**.** For every $n \in \mathbb{N}$, let $(x,y) \in \{0,1\}^n \times \{0,1\}^n$. Then

$$IP(x,y) = \bigoplus_{i=1}^{n} (x_i \wedge y_i)\,.$$

Both *Disj* and *IP* are *total bipartite Boolean functions*—that is, their input sets are bipartite and the function values are defined for every possible input pair.

Communication models. The study of communication complexity was initiated by Abelson [2] in the regime of real-valued messages and adapted by Yao [17] to the discrete regime that we are interested in. The models *P* and *BPP* that capture one's intuition of *efficient communication* (respectively, deterministic and randomised) date back to [17]. Later Babai, Frankl and Simon [3] introduced a number of stronger communication models—in particular, *AM* and *MA*—that intuitively corresponded to some classes studied in the context of structural computational complexity.

Definition 2.3 (*Polylogarithmic, P*)**.** We call deterministic bipartite communication protocols *P-protocols*.

We denote by *P* the class of functions solved by *P*-protocols of cost at most poly-log(n).

Definition 2.4 (*Bounded-error Probabilistic Polylogarithmic, BPP*)**.** For every $n \in \mathbb{N}$, let $f : \{0,1\}^n \times \{0,1\}^n \to \{0,1,\perp\}$ and $\varepsilon \geq 0$.

If for every input distribution μ_n there exists a *P*-protocol of cost at most $k_\varepsilon(n)$ that solves f with error at most ε, then we say that the BPP_ε-complexity of f, denoted by $BPP_\varepsilon(f)$, is at most $k_\varepsilon(n)$.

We let the *BPP*-complexity of f be its $BPP_{\frac{1}{3}}$-complexity.

We denote by *BPP* the class of functions whose *BPP*-complexity is at most poly-log(n).

The above definition of *BPP*, as well those among the following model definitions that are distribution-dependent, can be phrased in the "worst-case" formulations that do not make a reference to input distributions. Those variants usually correspond to the closures of our definitions with respect to mixed strategies, which, in turn, do not affect the resulting models, due to Von Neumann's minimax principle [16].

Definition 2.5 (*Non-deterministic Polylogarithmic, NP*)**.** For some $k(n)$, let $\Pi = \left\{ r_i \,\middle|\, i \in [2^{k(n)}] \right\}$ be a family of characteristic functions of combinatorial rectangles over $\{0,1\}^n \times \{0,1\}^n$.

We call such Π an *NP*-protocol of cost $k(n)$ that solves the function $f = \bigvee_{i=1}^{2^{k(n)}} r_i(x,y)$ (as well any partial g that is consistent with f on $g^{-1}(0) \cup g^{-1}(1)$).

We say that Π accepts every $f^{-1}(1)$ and rejects every $f^{-1}(0)$.

We say that Π solves a function $g : \{0,1\}^n \times \{0,1\}^n \to \{0,1,\perp\}$ with error ε with respect to an input distribution μ_n, if $\boldsymbol{Pr}_{(X,Y)\sim\mu_n}[f(X,Y) \neq g(X,Y)] = \varepsilon$.

We denote by *NP* the class of functions solved by *NP*-protocols of cost at most poly-log(n).

Definition 2.6 (*Arthur-Merlin, AM*)**.** For every $n \in \mathbb{N}$, let $f : \{0,1\}^n \times \{0,1\}^n \to \{0,1,\perp\}$.

If for every input distribution μ_n there exists an *NP*-protocol of cost at most $k(n)$ that solves f with error at most $\frac{1}{3}$, then we say that the *AM*-complexity of f, denoted by $AM(f)$, is at most $k(n)$.

We denote by *AM* the class of functions whose *AM*-complexity is at most poly-log(n).

As we mentioned already, *AM* is a very strong model of communication, for which we currently do not have any non-trivial lower bound. All the following classes can be viewed (and some of them have been defined) as "weaker forms" of *AM*: for all of them we already have strong lower bounds.

Definition 2.7 (*Merlin-Arthur, MA*)**.** For every $n \in \mathbb{N}$, let $f : \{0,1\}^n \times \{0,1\}^n \to \{0,1,\perp\}$.

If for some $k(n)$ there are functions $h_1, \ldots, h_{2^{k(n)}} : \{0,1\}^n \times \{0,1\}^n \to \{0,1,\perp\}$, whose *BPP*-complexity is at most $k(n)$, such that $f(x,y) \equiv \bigvee_{i=1}^{2^{k(n)}} h_i(x,y)$, then we say that the *MA*-complexity of f, denoted by $MA(f)$, is at most $k(n)$.

We call Merlin-Arthur (*MA*) the class of functions whose *MA*-complexity is at most poly-log(n).

Note that "$\bigvee$" of partial functions is defined as in (2.1).

Definition 2.8 (*Small-advantage Bounded-error Probabilistic Polylogarithmic, SBP*)**.** For every $n \in \mathbb{N}$, let $f : \{0,1\}^n \times \{0,1\}^n \to \{0,1,\perp\}$.

If for input distribution μ_n and some $\alpha > 0$ there exists a *P*-protocol Π of cost at most $k'(n)$ such that

$$\Pr_{(X,Y)\sim\mu_n}\left[\Pi \text{ accepts } (X,Y) \,\middle|\, f(X,Y)=1\right] \geq \alpha \quad \text{and}$$

$$\Pr_{(X,Y)\sim\mu_n}\left[\Pi \text{ accepts } (X,Y) \,\middle|\, f(X,Y)=0\right] \leq \frac{\alpha}{2},$$

then we call Π an *SBP*-protocol for f with respect to μ_n. The complexity of this protocol is $k'(n) + \log(1/\alpha)$ (note that the value of α may depend on both n and μ_n).

If with respect to every μ_n there exists a *SBP*-protocol for f of cost at most $k(n)$, then we say that the *SBP*-complexity of f, denoted by $SBP(f)$, is at most $k(n)$.

We denote by *SBP* the class of functions whose *SBP*-complexity is at most poly-log(n).

It was shown in [8, 4] that $MA \subseteq SBP \subseteq AM$, in [14] that $SBP \neq AM$ and in [9] that $SBP \neq MA$. Therefore,

$$MA \subset SBP \subset AM\,.\,^{4}$$

The following complexity measure is a core methodological notion for this work.

Definition 2.9 (*Layer complexity*)**.** Let Π be an *NP*-protocol for solving $f : \{0,1\}^n \times \{0,1\}^n \to \{0,1,\perp\}$, possibly with error.

We say that the protocol Π

- has layer complexity l if every $(x,y) \in \{0,1\}^n \times \{0,1\}^n$ belongs to at most l rectangles of Π;
- has 0-layer complexity l_0 if every $(x,y) \in f^{-1}(0)$ belongs to at most l_0 rectangles of Π;
- has 1-layer complexity l_1 if every $(x,y) \in f^{-1}(1)$ belongs to at most l_1 rectangles of Π.

[4] Unless stated otherwise, we implicitly assume *partial functions* as the default type of communication problems. In those cases when the object under consideration is a total function and the fact is significant for the context, that will be mentioned explicitly.

The concept of layer complexity in the context of nondeterministic communication is very natural and not new, dating back at least to [12] by Karchmer, Newman, Saks and Wigderson. We will use it extensively in order to analyse some previously known subclasses of *AM* with strong lower bounds and to define some new ones.

The following two classes were introduced quite recently by Göös, Pitassi and Watson [10].

Definition 2.10 (*Unambiguous Arthur-Merlin, UAM*)**.** For every $n \in \mathbb{N}$, let $f : \{0,1\}^n \times \{0,1\}^n \to \{0,1,\perp\}$.

If for some constant $\varepsilon < \frac{1}{2}$ and every input distribution μ_n there exists an *NP*-protocol of cost at most $k(n)$ and 1-layer complexity 1 that solves f with error at most ε, then we say that the *UAM*-complexity of f, denoted by $UAM(f)$, is at most $k(n)$.

We denote by *UAM* the class of functions whose *UAM*-complexity is at most poly-log(n).

Definition 2.11 (*Unambiguous Arthur-Merlin with perfect completeness, UAM_{compl}*)**.** For every $n \in \mathbb{N}$, let $f : \{0,1\}^n \times \{0,1\}^n \to \{0,1,\perp\}$.

If for every input distribution μ_n there exists an *NP*-protocol of cost at most $k(n)$ and 1-layer complexity 1 that solves f with perfect completeness (that is, it accepts every $(x,y) \in f^{-1}(1)$) and soundness error at most $\frac{1}{2}$ (that is, $\boldsymbol{Pr}_{\mu_n}\left[(X,Y) \text{ is accepted} \middle| f(X,Y) = 0\right] \leq \frac{1}{2}$), then we say that the UAM_{compl}-complexity of f, denoted by $UAM_{\text{compl}}(f)$, is at most $k(n)$.

We denote by UAM_{compl} the class of functions whose UAM_{compl}-complexity is at most poly-log(n).

The classes *AM*, *MA*, UAM_{compl} and *UAM* can be defined in an alternative, more "narrative" way, where an almighty *prover Merlin* interacts with a limited *verifier Arthur* (who, in turn, is a two-headed union of the *players Alice* and *Bob*). In the cases of *AM*, UAM_{compl} and *UAM* these variants correspond to the closures with respect to mixed strategies (that are equivalent to our definitions, as mentioned earlier).

Note that the error parameter in the definitions of *AM* and UAM_{compl} are fixed without loss of generality, while for *UAM* it may be any constant $\varepsilon < \frac{1}{2}$. In the first two cases the error can be trivially reduced to an arbitrary constant by repeating the protocol constant number of times; on the other hand, in the case of *UAM* the possibility of efficient error reduction is not known, so fixing a specific ε might result in weakening the model.[5]

It was shown in [10] that $NP \not\subseteq UAM$. They also showed that $UAM \not\subseteq SBP$ held in the context of *query complexity*, later in [9] this separation was generalised to the case of communication complexity, thus implying that *UAM* and *SBP* are incomparable:

$$UAM \not\subseteq SBP \text{ and } SBP \not\subseteq UAM\,.$$

On the other hand, *UAM* and *SBP* are the strongest previously known communication complexity classes contained in *AM* with non-trivial lower bounds, which makes it interesting to look for their "natural merge" and try to prove good lower bounds there. That will be the quest of the next section.

[5] In order to reduce the error via repetition, the answer of the new protocol should be the majority vote of the individual answers in the case of *AM*, and their logical conjunction in the case of UAM_{compl}. The problem with this approach for *UAM* stems from the fact that in order to perform *two-sided* error reduction via repetition, one must take the *majority vote* of the individual answers, which would ruin the required uniqueness of 1-certificates.

3 *Layered Arthur-Merlin*: getting as close to *AM* as we can

Let us try to construct as strong a communication model "under *AM*" as we can analyse.

We start by considering several slightly stronger modifications of *MA* that will emphasise the intuition behind the main definitions that will follow.

Definition 3.1 (*MA*′)**.** For every $n \in \mathbb{N}$, let $f : \{0,1\}^n \times \{0,1\}^n \to \{0,1,\perp\}$.

If for some $k(n)$ and $1 \le t(n) \le 2^{k(n)}$ there are functions $h_1, \ldots, h_{t(n)} : \{0,1\}^n \times \{0,1\}^n \to \{0,1,\perp\}$, whose $BPP_{1/3\cdot t^2(n)}$-complexity is at most $k(n)$, such that $f(x,y) \equiv \bigvee_{i=1}^{t(n)} h_i(x,y)$, then we say that the MA'-complexity of f, denoted by $MA'(f)$, is at most $k(n)$.

We call such $\{h_i \mid i \in [t]\}$ an MA'-protocol for f. We address the value t as the layer complexity of this protocol.[6]

Observe that

$$MA(f) \le MA'(f) \in O\left((MA(f))^2\right)$$

always holds: the inequality follows trivially from the definitions, and the containment results from the well-known fact that for every function h and $\varepsilon > 0$, $BPP_\varepsilon(h) \in O\left(BPP(h) \cdot \log \frac{1}{\varepsilon}\right)$. So, *MA* is the class of functions, whose MA'-complexity is at most poly-$\log(n)$.

Now suppose $MA'(f) \le k(n)$, what does it imply with respect to a *known* input distribution μ? In this case for every h_i there is a P-protocol of cost at most $k(n)$ that computes a function g_i, such that $\boldsymbol{Pr}_\mu[h_i(X,Y) \ne g_i(X,Y)] \le \frac{1}{3\cdot t^2(n)}$; accordingly, the union bound gives

$$\Pr_\mu\left[f(X,Y) \ne \bigvee_{i=1}^{t(n)} g_i(X,Y)\right] \le \frac{1}{3\cdot t(n)}.$$

What can we say about a communication complexity class that only requires that the above holds for every μ: in particular, what will be its relation to *MA*? Let us define it.

Definition 3.2 ($\overline{MA}$)**.** For every $n \in \mathbb{N}$, let $f : \{0,1\}^n \times \{0,1\}^n \to \{0,1,\perp\}$.

For some $k(n)$ and $1 \le t(n) \le 2^{k(n)}$, let $\Pi = \{g_i \mid i \in [t(n)]\}$ be a family of functions $g_i : \{0,1\}^n \times \{0,1\}^n \to \{0,1,\perp\}$, whose P-complexity is at most $k(n)$, such that for some input distribution μ_n it holds that $\boldsymbol{Pr}_{\mu_n}\left[f(X,Y) \ne \bigvee_{i=1}^{t(n)} g_i(X,Y)\right] \le 1/3t(n)$. Then we call Π an $\overline{MA}$-protocol of cost $k(n)$ for f with respect to μ_n.

If for every input distribution μ_n there exists an $\overline{MA}$-protocol of cost $k(n)$ for f, then we say that the $\overline{MA}$-complexity of f, denoted by $\overline{MA}(f)$, is at most $k(n)$.

We denote by $\overline{MA}$ the class of functions whose $\overline{MA}$-complexity is at most poly-$\log(n)$.

[6] Note the slight inconsistency in our use of the term *layer complexity* and the parameter t that denotes it: most of the time they stand for the actual number of rectangles that an input value belongs to, but occasionally—in particular, in the context of *MA*′—we use them for an *upper bound* on the number of possible "witnesses." The main reason for that is the lack of natural correspondence between the *BPP*-protocols (intrinsic to the definition of *MA*′) and combinatorial rectangles.

Note that

$$MA \subseteq \overline{MA}$$

follows from the definition and the previous discussion: $\overline{MA}(f) \leq MA'(f) \in O\left((MA(f))^2\right)$.

Towards our goal to construct a communication model under *AM* as strong as we can analyse, let us look at UAM_{compl}: together with *MA* these are, arguably, the two most natural (though not the strongest) "sub-*AM*" models for which we have good lower bounds. Conceptually, the insightful lower bounds given by Klauck [13] for *MA* and by Göös, Pitassi and Watson [10] for UAM_{compl} can be viewed as two different approaches to analysing strong "sub-*AM*" models of communication complexity.

On the one hand, the more recently defined UAM_{compl} has at least one important "*AM*-like" property that *MA* lacks: *AM* puts no limitations on the layer complexity of protocols; *MA* limits the number of "layers" over any input pair; UAM_{compl} and *UAM* only limit the 1-layer complexity (that is, they let the 0-layer complexity of a protocol be arbitrary, like *AM* and unlike *MA*). This difference seems to be rather crucial:

- While the lower-bound argument of [13] against *MA* can be generalised to work against a communication model that would limit only the 0-layer complexity of a protocol, it does not seem to go through if only the 1-layer complexity is limited.

- If we consider the natural (and the most common) situation when the target function is balanced with respect to its "hard" distribution—which is the case, for instance, for all functions with low discrepancy—then the "expected density" of protocol's rectangles over the points in the (erroneously) accepted ε-fraction of $f^{-1}(0)$ would be much higher than the density in the (rightly) accepted majority of $f^{-1}(1)$. In other words, the expected number of protocol's rectangles that an *accepted* $(x,y) \in f^{-1}(0)$ belongs to would be considerably higher than the analogous value for $(x,y) \in f^{-1}(1)$. Accordingly, limiting only the 1-layer complexity feels like a weaker restriction (i. e., resulting in a stronger defined model) than limiting only the 0-layer complexity (or both).

On the other hand, even though the classes[7] UAM_{compl} and *UAM* limit only the 1-layer complexity of a protocol, the actual quantitative limitation that they put is way too strong: it is 1, as opposed to the quasi-polynomial limitation on the (total) layer complexity of *MA* (as emphasised by Definition 3.2). For instance, it has been shown in [10] that $NP \not\subseteq UAM$ (note that $NP \subseteq MA$ and $UAM_{\text{compl}} \subseteq UAM$). To include *NP*, an "*NP*-like" class must allow protocols with super-constant 1-layer complexity.

On the technical level, comparing the definitions of $\overline{MA}$ (Definition 3.2) and of UAM_{compl} (Definition 2.11), we can see that in both cases the membership of a function f implies existence of a family of rectangles, whose union approximates f—that is, existence of good *NP*-approximations of f:

- if $f \in \overline{MA}$ (in particular, if $f \in MA$), then for some $t(n) \in \mathbb{N}$ and every input distribution μ_n there exists an *NP*-protocol of cost at most poly-log(n) and layer complexity at most $t(n)$ that solves f with error at most $1/3t(n)$;

[7] Most of the time the distinction between the two classes is insignificant for the context of this work. Nevertheless, we will rather often explicitly mention both *UAM* and UAM_{compl} in the same context, as, in our opinion, the former is more naturally defined, while the latter is somewhat simpler to analyse and often allows for clearer intuition.

- if $f \in UAM_{\text{compl}}$, then for every input distribution μ_n there exists an *NP*-protocol of cost at most poly-log(n) and 1-layer complexity 1 that solves f with perfect completeness and soundness error at most $1/2$ with respect to μ_n.

Note that the above membership condition of UAM_{compl} is sufficient, and that of $\overline{MA}$ is just necessary.

Let us use this intuition to define a new communication complexity class that includes both UAM_{compl} and $\overline{MA}$.

Definition 3.3 (*Layered Arthur-Merlin, LAM*)**.** For every $n \in \mathbb{N}$, let $f : \{0,1\}^n \times \{0,1\}^n \to \{0,1,\perp\}$.

If for input distribution μ_n there exists an *NP*-protocol Π of 1-layer complexity t that solves f with completeness error at most $1/3$ and soundness error at most $1/3t$, then we call Π a *LAM*-protocol for f with respect to μ_n. If Π contains K rectangles, then its complexity is $\log(K)$.

If with respect to every μ_n there exists a *LAM*-protocol for f of cost at most $k(n)$, then we say that the *LAM*-complexity of f, denoted by $LAM(f)$, is at most $k(n)$.

We denote by *LAM* the class of functions whose *LAM*-complexity is at most poly-log(n).

It follows readily from the previous discussion that

$$NP, MA, \overline{MA}, UAM_{\text{compl}} \subseteq LAM \subseteq AM.$$

To make it somewhat stronger and to simplify its definition, we have granted to *LAM* a few additional relaxations (not needed in order to include *MA* and UAM_{compl}): Most significantly, in *LAM* the layer complexity bound t can be chosen per distribution μ_n, and it does not have to be error-independent, unlike in the cases of both *MA* and UAM_{compl} (for the latter it equals 1).

Let us further strengthen the model, so that the corresponding complexity class would include all previously known subclasses of *AM* with strong lower bounds. The following definition can be viewed as *LAM* with relaxed accuracy requirements.

Definition 3.4 (*Small-advantage Layered Arthur-Merlin, SLAM*)**.** For every $n \in \mathbb{N}$, let $f : \{0,1\}^n \times \{0,1\}^n \to \{0,1,\perp\}$.

If for input distribution μ_n and some $\alpha > 0$ there exists an *NP*-protocol Π of 1-layer complexity t such that

$$\Pr_{(X,Y)\sim\mu_n}\left[\Pi \text{ accepts } (X,Y) \,\middle|\, f(X,Y)=1\right] \geq \alpha \text{ and}$$
$$\Pr_{(X,Y)\sim\mu_n}\left[\Pi \text{ accepts } (X,Y) \,\middle|\, f(X,Y)=0\right] \leq \frac{\alpha}{2t},$$

then we call Π a *SLAM*-protocol for f with respect to μ_n. If Π contains K rectangles, then its complexity is $\log(K/\alpha)$ (the value of α may depend on both n and μ_n).

If with respect to every μ_n there exists a *SLAM*-protocol for f of cost at most $k(n)$, then we say that the *SLAM*-complexity of f, denoted by $SLAM(f)$, is at most $k(n)$.

We denote by *SLAM* the class of functions whose *SLAM*-complexity is at most poly-log(n).

As any *LAM*-protocol of cost k is also a *SLAM*-protocol of cost $k + \log\frac{3}{2}$,

$$SLAM(f) < LAM(f) + 1$$

holds for all f.

Later (Section 3.2) we will see that *SLAM* indeed is a proper subclass of *AM* that includes all previously known (as far as we are aware) subclasses of *AM* with strong lower bounds:

$$NP, MA, \overline{MA}, UAM_{\text{compl}}, LAM, UAM, SBP \subseteq SLAM \subset AM;$$

moreover, it is strictly stronger than their union:

$$UAM \cup SBP \subset SLAM.$$

3.1 Limitations of *LAM* and *SLAM*

Let us see that the *SLAM*-complexity is a subject to the discrepancy bound.

Definition 3.5 (*Discrepancy*)**.** For every $n \in \mathbb{N}$, let $f : \{0,1\}^n \times \{0,1\}^n \to \{0,1,\perp\}$ and μ_n be a distribution on $\{0,1\}^n \times \{0,1\}^n$.

The discrepancy of f with respect to μ_n is defined as

$$\operatorname{disc}_{\mu_n}(f) = \max\left\{r \,\middle|\, \left|\mu_n(r \cap f^{-1}(1)) - \mu_n(r \cap f^{-1}(0))\right|\right\},$$

where r ranges over the combinatorial rectangles over $\{0,1\}^n \times \{0,1\}^n$.

We denote $\operatorname{disc}(f) = \min\left\{\mu \,\middle|\, \operatorname{disc}_\mu(f)\right\}$.

Theorem 3.6. *For any $f : \{0,1\}^n \times \{0,1\}^n \to \{0,1,\perp\}$:*

$$SLAM(f) \in \Omega\left(\sqrt{\log \frac{1}{\operatorname{disc}(f)}}\right).$$

That is,

$$SLAM \subseteq PP,$$

where *PP* is the class consisting of functions with high discrepancy. Along with other mentioned properties, this implies

$$SBP \cup UAM \subset SLAM \subseteq AM \cap PP \text{ and } SLAM \subset AM,$$

as $AM \not\subseteq PP$ is known [14].

Corollary 3.7. *Any LAM- or SLAM-protocol for IP has cost $\Omega(\sqrt{n})$.*

To prove the theorem we will use the following combinatorial lemma.

Lemma 3.8. *Let $C_1, \ldots, C_m$ be finite sets and $W \stackrel{\text{def}}{=} \bigcup_{i=1}^m C_i$. Let $t \in \mathbb{N}$, $\beta > 1$ and*

$$W_0 \stackrel{\text{def}}{=} \left\{w \in W \,\middle|\, 1 \le \left|\left\{i \in [m] \,\middle|\, w \in C_i\right\}\right| \le t\right\},$$
$$W_1 \stackrel{\text{def}}{=} \left\{w \in W \,\middle|\, \left|\left\{i \in [m] \,\middle|\, w \in C_i\right\}\right| \ge \beta \cdot t\right\}.$$

Let μ be a distribution on W, such that

$$\frac{\mu(W_1)}{\mu(W_0)} \geq \lambda$$

for some $\lambda > 0$. [8]

Then for any $\gamma > \lambda$ there exists $J \subseteq [m]$ of size at most $k \stackrel{\text{def}}{=} \left\lceil \log_{\frac{\beta+1}{2}} \left(\frac{\gamma}{\lambda}\right) \right\rceil$, such that for $C_J \stackrel{\text{def}}{=} \bigcap_{j\in J} C_j$ it holds that

$$\frac{\mu(C_J \cap W_1)}{\mu(C_J \cap W_0)} \geq \gamma, \text{ [8]} \tag{3.1}$$

and

$$\mu(C_J \cap W_1) \geq \mu(W_1) \cdot \left(\frac{\min\{1, \beta - 1\}}{2m}\right)^k.$$

Informally, the lemma says that for any family of sets $C_1, \ldots, C_m$ there exists $C_J = \bigcap_{j\in J} C_j$ that "highlights" points that are contained in more than a certain threshold number of sets C_i; moreover, the fraction of such points in W that end up in $C_J \cap W$ is not too small.

Proof. We will find a set C_{i_0} such that $\mu(C_{i_0} \cap W_1)$ is not too small and, at the same time, the ratio $\mu(C_{i_0} \cap W_1)/\mu(C_{i_0} \cap W_0)$ is significantly larger than $\mu(W_1)/\mu(W_0)$. The result will follow by induction on t.

The first part of the argument is the same for the base case ($t = 1$) and the inductive step ($t \geq 2$). Let $C_i(\cdot)$ denote the characteristic function of C_i, then

$$\begin{aligned}
&\mathop{\boldsymbol{E}}_{X\sim\mu}\left[\sum_{i=1}^m C_i(X) \middle| X \in W_1\right] \geq \beta \cdot \mathop{\boldsymbol{E}}_{X\sim\mu}\left[\sum_{i=1}^m C_i(X) \middle| X \in W_0\right];\\
&\sum_{i=1}^m \mathop{\boldsymbol{E}}_{X\sim\mu}\left[C_i(X) \middle| X \in W_1\right] \geq \beta \cdot \sum_{i=1}^m \mathop{\boldsymbol{E}}_{X\sim\mu}\left[C_i(X) \middle| X \in W_0\right];\\
&\sum_{i=1}^m \left(\mathop{\boldsymbol{E}}_{X\sim\mu}\left[C_i(X) \middle| X \in W_1\right] - \beta \cdot \mathop{\boldsymbol{E}}_{X\sim\mu}\left[C_i(X) \middle| X \in W_0\right]\right) \geq 0;\\
0 \;\leq\; &\sum_{i=1}^m \left(\frac{\beta+1}{2\beta} \cdot \mathop{\boldsymbol{E}}_{X\sim\mu}\left[C_i(X) \middle| X \in W_1\right] - \frac{\beta+1}{2} \cdot \mathop{\boldsymbol{E}}_{X\sim\mu}\left[C_i(X) \middle| X \in W_0\right]\right)\\
&= \sum_{i=1}^m \left(\mathop{\boldsymbol{E}}_{X\sim\mu}\left[C_i(X) \middle| X \in W_1\right] - \frac{\beta-1}{2\beta} \cdot \mathop{\boldsymbol{E}}_{X\sim\mu}\left[C_i(X) \middle| X \in W_1\right]\right.\\
&\qquad \left. - \frac{\beta+1}{2} \cdot \mathop{\boldsymbol{E}}_{X\sim\mu}\left[C_i(X) \middle| X \in W_0\right]\right);\\
&\sum_{i=1}^m \left(\mathop{\boldsymbol{E}}_{X\sim\mu}\left[C_i(X) \middle| X \in W_1\right] - \frac{\beta+1}{2} \cdot \mathop{\boldsymbol{E}}_{X\sim\mu}\left[C_i(X) \middle| X \in W_0\right]\right)
\end{aligned}$$

[8] Here let $\frac{x}{0} > y$ hold for any $x, y > 0$.

$$\geq \frac{\beta-1}{2\beta}\cdot \mathop{\boldsymbol{E}}_{X\sim\mu}\left[\sum_{i=1}^{m} C_i(X)\middle| X\in W_1\right]$$
$$\geq \frac{\beta-1}{2\beta}\cdot\beta\cdot t \;\geq\; \frac{\beta-1}{2};$$
$$\exists i_0\in[m]:\;\; \mathop{\boldsymbol{E}}_{X\sim\mu}\left[C_{i_0}(X)\middle| X\in W_1\right]-\frac{\beta+1}{2}\cdot\mathop{\boldsymbol{E}}_{X\sim\mu}\left[C_{i_0}(X)\middle| X\in W_0\right] \;\geq\; \frac{\beta-1}{2m}.$$

That is,

$$\frac{\mu(C_{i_0}\cap W_1)}{\mu(W_1)}-\frac{\beta+1}{2}\cdot\frac{\mu(C_{i_0}\cap W_0)}{\mu(W_0)}\geq\frac{\beta-1}{2m},$$

which implies

$$\mu(C_{i_0}\cap W_1)\geq\frac{\beta-1}{2m}\cdot\mu(W_1) \tag{3.2}$$

and

$$\frac{\mu(C_{i_0}\cap W_1)}{\mu(C_{i_0}\cap W_0)}\geq\frac{\beta+1}{2}\cdot\frac{\mu(W_1)}{\mu(W_0)}\geq\frac{\beta+1}{2}\cdot\lambda. \tag{3.3}$$

At this point we check whether letting $J=\{i_0\}$ would satisfy the statement of the lemma. Assume that it would not; as $\gamma>\lambda\Rightarrow k\geq 1$, this necessarily means that (3.3) is insufficient to guarantee (3.1). In other words, it holds that

$$\frac{\beta+1}{2}\cdot\lambda\leq\frac{\mu(C_{i_0}\cap W_1)}{\mu(C_{i_0}\cap W_0)}<\gamma,$$

where the latter inequality is the contrapositive of (3.1) with respect to $J=\{i_0\}$, and therefore

$$\frac{\beta+1}{2}<\frac{\gamma}{\lambda}. \tag{3.4}$$

Denote

$$C'_j\stackrel{\text{def}}{=}C_j\cap C_{i_0}$$

for all $j\neq i_0$ and

$$C''_{i_0}\stackrel{\text{def}}{=}C_{i_0}\setminus\bigcup_{j\neq i_0}C_j.$$

Note that $C''_{i_0}\subseteq W_0$.

How we continue from here depends on the value of t: first suppose that $t=1$ (the base case for the induction). Let $i_1\in[m]\setminus\{i_0\}$ be such that $\mu(C'_{i_1}\cap W_1)$ is maximised. Let $J\stackrel{\text{def}}{=}\{i_0,i_1\}$. From (3.4) it follows that

$$k=\left\lceil\log_{\frac{\beta+1}{2}}\left(\frac{\gamma}{\lambda}\right)\right\rceil\geq 2=|J|.$$

From (3.2) and the choice of i_1,

$$\mu(C_J \cap W_1) \geq \frac{1}{m-1} \cdot \frac{\beta-1}{2m} \cdot \mu(W_1) > \mu(W_1) \cdot \left(\frac{\min\{1,\beta-1\}}{2m}\right)^2 .$$

As $t = 1$,

$$\mu(C_J \cap W_0) = 0 ,$$

which satisfies (3.1) (according to Footnote 8). This finishes the proof of the base case.

Now suppose $t \geq 2$. That is, we are inside the inductive step, so let us apply Lemma 3.8 inductively to the family $\left\{C'_j \,\middle|\, j \neq i_0\right\}$ with parameters

$$m' = m-1,\ t' = t-1,\ \beta' = \frac{\beta \cdot t - 1}{t-1},\ \gamma' = \gamma.$$

Note that this choice corresponds to

$$W' = (W \cap C_{i_0}) \setminus C''_{i_0},\ W'_0 = (W_0 \cap C_{i_0}) \setminus C''_{i_0},\ W'_1 = W_1 \cap C_{i_0} ,$$
$$\lambda' = \frac{\beta+1}{2} \cdot \lambda ,$$

where the last equality follows from (3.3). Note that $\lambda' < \gamma'$ follows from (3.4).

The lemma guarantees existence of (non-empty) $J' \subseteq [m] \setminus \{i_0\}$ of size at most

$$k' = \left\lceil \log_{\frac{\beta'+1}{2}} \left(\frac{\gamma}{\lambda'}\right) \right\rceil \leq \left\lceil \log_{\frac{\beta+1}{2}} \left(\frac{\gamma}{\lambda'}\right) \right\rceil = \left\lceil \log_{\frac{\beta+1}{2}} \left(\frac{\gamma}{\lambda}\right) - 1 \right\rceil = k-1$$

(where the inequality follows from $\beta' > \beta$), such that for $C_{J'} \stackrel{\text{def}}{=} \bigcap_{j \in J'} C'_j$ it holds that

$$\frac{\mu(C_{J'} \cap C_{i_0} \cap W_1)}{\mu(C_{J'} \cap C_{i_0} \cap W_0)} \geq \frac{\mu(C_{J'} \cap W'_1)}{\mu(C_{J'} \cap W'_0)} \geq \gamma;$$
$$\mu(C_{J'} \cap C_{i_0} \cap W_1) = \mu(C_{J'} \cap W'_1) \geq \mu(W'_1) \cdot \left(\frac{\min\{1,\beta-1\}}{2m}\right)^{k'}$$
$$\geq \mu(C_{i_0} \cap W_1) \cdot \left(\frac{\min\{1,\beta-1\}}{2m}\right)^{k-1} \geq \mu(W_1) \cdot \left(\frac{\min\{1,\beta-1\}}{2m}\right)^{k},$$

where the last inequality follows from (3.2). Letting $J \stackrel{\text{def}}{=} J' \cup \{i_0\}$ finishes the proof. □

We are ready to prove the lower bound.

Proof. The argument is as follows. Recall that the core advantage of the models *LAM* and *SLAM* over *MA* is allowing arbitrarily high 0-layer complexity in efficient protocols: If the 0-layer complexity of a given protocol Π is not above its 1-layer complexity, then Klauck's argument for *MA* limits Π's strength.

The case when the 0-layer complexity is high was the main challenge of this work and the reason why it was written. Here we further assume that the *average* 0-layer complexity of Π is *noticeably* higher than its *average* 1-layer complexity (we notice that the other cases can be handled by a straightforward adaptation of Klauck's argument). The layer complexity measures the density of Π's rectangles, and the assumed difference in the average densities between $f^{-1}(0)$ and $f^{-1}(1)$ implies that the membership of a pair of random variables (X,Y) in "too many" rectangles makes the event $[f(X,Y)=0]$ more likely. Lemma 3.8 gives us a "not too small" rectangle intersection—therefore a rectangle by itself—where many elements belong to many (original) rectangles. The discrepancy assumption applied to this new rectangle concludes the argument.

Let μ be a distribution that achieves $\mathrm{disc}(f)=\mathrm{disc}_\mu(f)$ and

$$\Pi=\left\{r_i\,\middle|\, i\in[2^{k(n)}]\right\}$$

be a *SLAM*-protocol for f with respect to μ of cost $k(n)+\log(1/\alpha(n))$ that accepts $\alpha(n)$-fraction of the elements of $f^{-1}(1)$, and whose 1-layer complexity is $t(n)$.

By the definition of *SLAM*, the soundness error of Π in solving f with respect to μ is at most

$$\frac{\alpha(n)}{2\cdot t(n)}. \tag{3.5}$$

By the definition of disc_μ (and the fact that $\{0,1\}^n\times\{0,1\}^n$ is a rectangle),

$$\underset{(X,Y)\sim\mu}{\boldsymbol{Pr}}[f(X,Y)=0],\ \underset{(X,Y)\sim\mu}{\boldsymbol{Pr}}[f(X,Y)=1]\ \in\ \frac{1}{2}\pm\frac{\mathrm{disc}_\mu(f)}{2}=\frac{1}{2}\pm\frac{\mathrm{disc}(f)}{2}. \tag{3.6}$$

Let $l_0^{av}(n)$ denote the *average* 0-layer complexity of Π, namely

$$l_0^{av}(n)\stackrel{\mathrm{def}}{=}\underset{(X,Y)\sim\mu}{\boldsymbol{E}}\left[\left|\left\{r\in\Pi\,\middle|\,(X,Y)\in r\right\}\right|\,\middle|\,(X,Y)\in\Pi^{-1}(1)\cap f^{-1}(0)\right],$$

where "$\Pi^{-1}(1)$" denotes the set of input pairs accepted by Π.

Fix $n\in\mathbb{N}$. We will consider two cases, distinguished by the value of $l_0^{av}(n)$.

First suppose that

$$l_0^{av}(n)\le\frac{4\cdot t(n)}{3}. \tag{3.7}$$

Then

$$\begin{aligned}\sum_{r\in\Pi}\mu(r\cap f^{-1}(1))&\ge\underset{(X,Y)\sim\mu}{\boldsymbol{Pr}}\left[(X,Y)\in\Pi^{-1}(1)\cap f^{-1}(1)\right]\\&=\underset{\mu}{\boldsymbol{Pr}}[f(X,Y)=1]\cdot\underset{\mu}{\boldsymbol{Pr}}\left[(X,Y)\in\Pi^{-1}(1)\middle|f(X,Y)=1\right]\\&\ge\left(\frac{1}{2}-\frac{\mathrm{disc}(f)}{2}\right)\cdot\alpha(n),\end{aligned}$$

where the last inequality follows from (3.6), and

$$\begin{aligned}
\sum_{r\in\Pi}\mu(r\cap f^{-1}(0)) &= \sum_{r\in\Pi}\sum_{(x,y)\in r\cap f^{-1}(0)}\mu(x,y)\\
&= \sum_{(x,y)\in f^{-1}(0)}\mu(x,y)\cdot\left|\left\{r\in\Pi\,\middle|\,(x,y)\in r\right\}\right|\\
&= \Pr_{(X,Y)\sim\mu}\left[f(X,Y)=0\right]\cdot\mathop{\boldsymbol{E}}_{(X,Y)\sim\mu}\left[\left|\left\{r\in\Pi\,\middle|\,(X,Y)\in r\right\}\right|\middle|f(X,Y)=0\right]\\
&= \Pr_{\mu}\left[(X,Y)\in\Pi^{-1}(1)\cap f^{-1}(0)\right]\\
&\qquad\cdot\mathop{\boldsymbol{E}}_{\mu}\left[\left|\left\{r\in\Pi\,\middle|\,(X,Y)\in r\right\}\right|\middle|(X,Y)\in\Pi^{-1}(1)\cap f^{-1}(0)\right]\\
&= \Pr_{\mu}\left[(X,Y)\in\Pi^{-1}(1)\cap f^{-1}(0)\right]\cdot l_0^{av}(n)\\
&\le \Pr_{\mu}\left[f(X,Y)=0\right]\cdot\frac{\alpha(n)}{2\cdot t(n)}\cdot\frac{4\cdot t(n)}{3}\\
&\le \left(\frac{1}{2}+\frac{\mathrm{disc}(f)}{2}\right)\cdot\frac{2\cdot\alpha(n)}{3},
\end{aligned}$$

where the first inequality follows from (3.5) and (3.7), and the last one from (3.6). Therefore,

$$\sum_{r\in\Pi}\mu(r\cap f^{-1}(1))-\mu(r\cap f^{-1}(0))\ge\left(\frac{1}{6}-\frac{5\cdot\mathrm{disc}(f)}{6}\right)\cdot\alpha(n)$$

and for some $r_0\in\Pi$ it holds that

$$\mathrm{disc}(f)\ge\mu(r_0\cap f^{-1}(1))-\mu(r_0\cap f^{-1}(0))\ge\left(\frac{1}{6}-\frac{5\cdot\mathrm{disc}(f)}{6}\right)\cdot\alpha(n)\cdot 2^{-k(n)}$$

and

$$k(n)+\log\left(\frac{1}{\alpha(n)}\right)\in\log\left(\frac{1}{\mathrm{disc}(f)}\right)-O(1),$$

as required.

Now suppose that

$$l_0^{av}(n)>\frac{4\cdot t(n)}{3}.$$

Define

$$A\stackrel{\text{def}}{=}\left\{(x,y)\,\middle|\,\left|\left\{r\in\Pi\,\middle|\,(x,y)\in r\right\}\right|\ge\frac{5\cdot t(n)}{4}\right\}.$$

Let us see that $\mu(A)$ cannot be too small.

$$\begin{aligned}\frac{4\cdot t(n)}{3} < l_0^{av}(n) &= \mathop{\boldsymbol{E}}_{\mu}\left[\left|\left\{r\in\Pi\,\middle|\,(X,Y)\in r\right\}\right|\middle|(X,Y)\in\Pi^{-1}(1)\cap f^{-1}(0)\right]\\ &\le \mathop{\boldsymbol{Pr}}_{\mu}\left[(X,Y)\in A\middle|(X,Y)\in\Pi^{-1}(1)\cap f^{-1}(0)\right]\cdot 2^{k(n)}\\ &\quad+\left(1-\mathop{\boldsymbol{Pr}}_{\mu}\left[(X,Y)\in A\middle|(X,Y)\in\Pi^{-1}(1)\cap f^{-1}(0)\right]\right)\cdot\frac{5\cdot t(n)}{4}\\ &\le \frac{5\cdot t(n)}{4}+2^{k(n)}\cdot\mathop{\boldsymbol{Pr}}_{\mu}\left[(X,Y)\in A\middle|(X,Y)\in\Pi^{-1}(1)\cap f^{-1}(0)\right].\end{aligned}$$

Therefore,

$$\mathop{\boldsymbol{Pr}}_{\mu}\left[(X,Y)\in A\middle|(X,Y)\in\Pi^{-1}(1)\cap f^{-1}(0)\right] > \frac{t(n)\cdot 2^{-k(n)}}{12}$$

and

$$\mathop{\boldsymbol{Pr}}_{\mu}\left[(X,Y)\in A\right] > \frac{\mu\left(\Pi^{-1}(1)\cap f^{-1}(0)\right)}{12\cdot 2^{k(n)}}.$$

On the other hand,

$$\mu\left(\Pi^{-1}(1)\cap f^{-1}(0)\right)\ge\max\left\{r\in\Pi\middle|\mu\left(r\cap f^{-1}(0)\right)\right\}\ge\frac{\alpha(n)}{2^{k(n)}}\cdot\left(\frac{1}{2}-\frac{\mathrm{disc}(f)}{2}\right)^2,$$

where the last inequality follows from the fact that the μ-weight of the largest rectangle of Π is, due to (3.6), at least

$$\frac{\alpha(n)}{2^{k(n)}}\cdot\left(\frac{1}{2}-\frac{\mathrm{disc}(f)}{2}\right),$$

and the relative μ-weight of $f^{-1}(0)$ in it is at least

$$\frac{1}{2}-\frac{\mathrm{disc}(f)}{2}.$$

Assuming $\left[\mathrm{disc}(f)\le\frac{1}{2}\right]$ (otherwise the desired statement holds trivially), we get

$$\mu(A)\ge\frac{\alpha(n)}{192\cdot 2^{2\cdot k(n)}}. \tag{3.8}$$

Let

$$B\stackrel{\text{def}}{=}\left\{(x,y)\,\middle|\,1\le\left|\left\{r\in\Pi\,\middle|\,(x,y)\in r\right\}\right|\le t(n)\right\},$$

then

$$\mu(B)\le\mu\left(\Pi^{-1}(1)\right)<\alpha(n), \tag{3.9}$$

as Π accepts, with respect to μ, $\alpha(n)$-fraction of $f^{-1}(1)$ and smaller fraction of $f^{-1}(0)$.

Note that

$$A \subset f^{-1}(0) \text{ and } f^{-1}(1) \cap \Pi^{-1}(1) \subseteq B, \tag{3.10}$$

as follows from their definitions and the fact that the 1-layer complexity of Π is $t(n)$.

Let us make use of the difference in the "rectangle density" of A and B via applying Lemma 3.8. Namely, let $m \stackrel{\text{def}}{=} 2^{k(n)}$, $C_i \stackrel{\text{def}}{=} r_i$, $t \stackrel{\text{def}}{=} t(n)$, $\beta \stackrel{\text{def}}{=} \frac{5}{4}$ and $\gamma \stackrel{\text{def}}{=} 2$. Then the conditions of the lemma hold with respect to $W_0 = B$, $W_1 = A$ and $\lambda = \frac{1}{192 \cdot 2^{2 \cdot k(n)}}$. Then there exists some $J \subseteq [2^{k(n)}]$, such that for

$$s \stackrel{\text{def}}{=} \bigcap_{j \in J} r_j$$

it holds that

$$\frac{\mu(s \cap A)}{\mu(s \cap B)} \geq 2$$

and

$$\mu(s \cap A) \in \alpha(n) \cdot 2^{-O\left(k^2(n)\right)},$$

as follows from (3.8), (3.9) and the statement of the lemma.

That is,

$$\mu(s \cap A) - \mu(s \cap B) \geq \frac{\mu(s \cap A)}{2} \in \alpha(n) \cdot 2^{-O\left(k^2(n)\right)}.$$

As $s \subseteq \Pi^{-1}(1)$, it follows from (3.10) that

$$\mu(s \cap f^{-1}(0)) - \mu(s \cap f^{-1}(1)) \in \alpha(n) \cdot 2^{-O\left(k^2(n)\right)},$$

and since s is a rectangles' intersection and therefore a rectangle itself,

$$\operatorname{disc}(f) \in \alpha(n) \cdot 2^{-O\left(k^2(n)\right)},$$

that is,

$$O\left(k^2(n)\right) + \log\left(\frac{1}{\alpha(n)}\right) \geq \log\left(\frac{1}{\operatorname{disc}(f)}\right) \quad \Rightarrow \quad k(n) + \log\left(\frac{1}{\alpha(n)}\right) \in \Omega\left(\sqrt{\log \frac{1}{\operatorname{disc}(f)}}\right),$$

as required. □

3.2 More about *LAM* and *SLAM*

When we defined the communication complexity class *SLAM* (Definition 3.4), we promised to show later that

$$NP, MA, UAM_{\text{compl}}, UAM, SBP, LAM \subseteq SLAM \subset AM$$

and

$$UAM \cup SBP \subset SLAM\,.$$

The relations

$$NP, MA, UAM_{\text{compl}} \subseteq LAM \subseteq AM \text{ and } LAM \subseteq SLAM$$

follow trivially from the definitions. Theorem 3.6 implies that

$$SLAM \subseteq PP \implies SLAM \subset AM\,,$$

as *PP* is the class consisting of functions with high discrepancy and $AM \nsubseteq PP$ is known [14].

It remains to see that

$$UAM \cup SBP \subset SLAM \subseteq AM\,,$$

which will be implied by the upcoming Claims 3.9 and 3.13.

Claim 3.9. *For any bipartite Boolean function f:*

$$AM(f) \in O\left(SLAM(f) + \log n\right).$$

Proof. The proof combines the "randomness sparsification" method of Goldwasser and Sipser [8] with *NP*-witnessing.

Assume $SLAM(f) = k(n)$. Then by Von Neumann's minimax principle [16] there exists a family $\Pi = \{h_1, \ldots, h_m\}$ for some $m \in \mathbb{N}$, where every h_i is a bipartite Boolean function computable by an *NP*-protocol of cost at most $k(n)$, such that

$$\forall (x,y): \ \left|\left\{i \in [m] \,\middle|\, h_i(x,y) = 1\right\}\right| \ \begin{cases} \geq \alpha \cdot m & \text{if } f(x,y) = 1\,, \\ \leq \frac{\alpha}{2} \cdot m & \text{if } f(x,y) = 0\,, \end{cases}$$

for some $\alpha \geq 2^{-k(n)}$.[9]

By a standard Bernstein-type concentration argument (e. g., [7], Lemma 1, *Hoeffding inequality*), there exists $l \in O\left(n/\alpha\right) \subseteq O\left(2^{k(n)+\log n}\right)$ such that for some $\Pi' \subseteq \Pi$ of size l it holds that

$$\forall (x,y): \ \left|\left\{i \in [l] \,\middle|\, h_i(x,y) = 1\right\}\right| \ \begin{cases} \geq \frac{2\alpha}{3} \cdot l & \text{if } f(x,y) = 1\,, \\ \leq \frac{\alpha}{3} \cdot l & \text{if } f(x,y) = 0\,, \end{cases}$$

[9] Note that the actual value of t is insignificant for this argument, which is not surprising: if we modify the definition of *SLAM* (Definition 3.4) by allowing arbitrary 1-layer complexity, then we end up with a definition of *AM*.

where we have assumed without loss of generality that $\Pi' = \{h_1, \ldots, h_l\}$.

By another application of the Hoeffding inequality, for some $s \in O\left(n + \frac{1}{\alpha}\right) \subseteq O\left(2^{k(n)+\log n}\right)$ and a uniformly random function $g : [l] \to [s]$ it holds with positive probability that

$$\forall (x,y) : \ \underset{Z \in [s]}{\boldsymbol{Pr}} \left[\exists i \in [l] : \ h_i(x,y) = 1 \wedge g(i) = Z\right] \begin{cases} \geq \frac{3}{5} & \text{if } f(x,y) = 1\,, \\ \leq \frac{2}{5} & \text{if } f(x,y) = 0\,. \end{cases} \tag{3.11}$$

Fix any such g.

Consider the following *AM*-protocol (described below in a distribution-free regime, which is the dual equivalent of Definition 2.6).

- The players pick $Z \in [s]$ and send it to Merlin.
- Merlin responds with $i \in [l]$ and $w \in \{0,1\}^{k(n)}$.
- The players accept if and only if $h_i(X,Y) = 1 \wedge g(i) = Z$, where the former is witnessed by w (recall that $NP(h_i) \leq k(n)$).

By (3.11), this is an *AM*-protocol for f with error at most $2/5$; repeating it several times and taking the majority vote brings the error bound to at most $1/3$. The cost of the resulting protocol is in $O(k(n) + \log n)$, as required. □

To see that $UAM \cup SBP \subset SLAM$, we prove a somewhat stronger separation:

$$LAM \nsubseteq UAM \cup SBP\,.$$

For that we will use several results from [10, 9].

Fact 3.10 (*$NP \nsubseteq UAM$ [10]*)**.** *Let $\neg Disj(x,y) \stackrel{\text{def}}{=} 1 - Disj(x,y)$ for every $(x,y) \in \{0,1\}^n \times \{0,1\}^n$. Then*

$$UAM(\neg Disj) \in \Omega(n)\,.$$

The following partial function has been used to show that $UAM_{\text{compl}} \nsubseteq SBP$.

Definition 3.11 (*Gut-IP_n [10, 9]*)**.** For an even $m \in \mathbb{N}$, let $n = m^2 \cdot \lceil 200 \cdot \log m \rceil$. For any $x \in \{0,1\}^n$ and $i, j \in [m]$, let $x_{i,j}$ denote the sub-string of x that starts from bit $\lceil 200 \cdot \log m \rceil \cdot (m \cdot (i-1) + j - 1) + 1$ and contains $\lceil 200 \cdot \log m \rceil$ bits. Denote

$$\forall i \in [m] : \ \sharp_i \stackrel{\text{def}}{=} \left|\left\{ j \,\middle|\, IP(x_{i,j}, y_{i,j}) = 1 \right\}\right|,$$

then

$$Gut\text{-}IP_n(x,y) = \begin{cases} 1 & \text{if } \forall i : \sharp_i = 1; \\ 0 & \text{if } \left|\left\{ i \,\middle|\, \sharp_i = 0 \right\}\right| = \left|\left\{ i \,\middle|\, \sharp_i = 2 \right\}\right| = \frac{m}{2}; \\ \perp & \text{otherwise.} \end{cases}$$

Fact 3.12 ($UAM_{\text{compl}} \not\subseteq SBP$ *[10, 9]*)**.**

$$UAM_{\text{compl}}(\textit{Gut-IP}_n) \in O(\log n);$$
$$SBP(\textit{Gut-IP}_n) \in \Omega\left(\sqrt{m} \cdot \log m\right) = \Omega\left(n^{\frac{1}{4}} \cdot \log^{\frac{3}{4}} n\right).$$

Claim 3.13. *For any $n \in \mathbb{N}$ such that Gut-IP$_n$ is defined and $x_1, x_2, y_1, y_2 \in \{0,1\}^n$, let*

$$f((x_1,x_2),(y_1,y_2)) \stackrel{\text{def}}{=} \neg \textit{Disj}(x_1,y_1) \wedge \textit{Gut-IP}_n(x_2,y_2).$$

Then

$$LAM(f), SLAM(f) \in O\left(\log^2 n\right);$$
$$UAM(f) \in \Omega(n);$$
$$SBP(f) \in \Omega\left(n^{\frac{1}{4}} \cdot \log^{\frac{3}{4}} n\right).$$

Proof. Consider an input distribution μ that fixes $X_1 = Y_1 = 1^n$ and makes the pair (X_2, Y_2) come from a hard distribution for $\textit{Gut-IP}_n(X_2, Y_2)$, then any *SBP*-protocol that solves f with respect to μ must have complexity $\Omega\left(n^{1/4} \cdot \log^{3/4} n\right)$, according to Fact 3.12. Similarly, a distribution that fixes $(X_2, Y_2) \in \textit{Gut-IP}_n^{-1}(1)$ arbitrarily and makes $\textit{Disj}(X_1, Y_1)$ hard for *UAM* witnesses that $UAM(f) \in \Omega(n)$, according to Fact 3.10.

To see that $LAM(f) \in O\left(\log^2 n\right)$, let μ be any input distribution for f and let μ' be the marginal distribution of (X_2, Y_2) when $((X_1, X_2), (Y_1, Y_2)) \sim \mu$. Consider a UAM_{compl}-protocol Π of complexity $O(\log n)$ that solves $\textit{Gut-IP}_n$ with perfect completeness and soundness error at most $\frac{1}{2}$ with respect to μ', and let Π' be its amplified version of complexity $O\left(\log^2 n\right)$ that solves $\textit{Gut-IP}_n$ with soundness error at most $\frac{1}{3n}$ with respect to μ'.

Let $\Pi''((X_1, X_2), (Y_1, Y_2))$ be a nondeterministic protocol for f that does the following:

- emulates the behaviour of $\Pi'(X_2, Y_2)$;
- runs the trivial *NP*-protocol for $\neg \textit{Disj}(X_1, Y_1)$;
- accepts if and only if the two steps above have accepted.

The complexity of Π'' is in $O\left(\log^2 n\right)$.

Since an *NP*-protocol for $\neg \textit{Disj}$ is exact (though nondeterministic), an error can come only from the first step; since Π' has perfect completeness, so does Π''. The soundness error of Π'' in solving f with respect to μ equals that of Π' in solving $\textit{Gut-IP}_n$ with respect to μ', which is at most $\frac{1}{3n}$. Since Π' has 1-layer complexity 1, the 1-layer complexity of Π'' equals that of the *NP*-protocol for $\neg \textit{Disj}$, which is n. So, Π'' is a valid *LAM*-protocol for f with respect to μ, as required. □

4 On proving super-$\sqrt{n}$ lower bounds against MA

When Klauck [13] showed that $MA(Disj), MA(IP) \in \Omega(\sqrt{n})$, many believed that the actual MA-complexity of these problems was in $\Omega(n)$. So, it came as a surprise when Aaronson and Wigderson [1] demonstrated MA-protocols for $Disj$ and IP of cost $O(\sqrt{n}\log n)$, later improved by Chen [6] to $O\left(\sqrt{n\log n\log\log n}\right)$. That highlighted the importance of proving the "ultimate" lower bound of $\Omega(n)$ for the MA-complexity of any explicit communication problem.

We will define a communication model $\widetilde{MA}$ (Definition 4.1) that can be viewed as "non-uniform MA." Non-uniformity is the only possible source of advantage of $\widetilde{MA}$ over MA: we will see (Claim 4.2) that imposing the "uniformity constraint" on $\widetilde{MA}$-protocols makes them not stronger than MA-protocols. All known lower bounds on $MA(f)$ readily translate to $\widetilde{MA}(f)$. Intuitively, a lower bound argument that explores the uniformity of MA (as opposed to $\widetilde{MA}$) must have a very unusual structure.

We will see (Theorem 4.3) that for any f it holds that $\widetilde{MA}(f) \in O\left(\sqrt{n \cdot AM(f)}\right)$; in other words, any lower bound of the form $\widetilde{MA}(f) \in \omega(\sqrt{n})$ will have non-trivial consequences for $AM(f)$.[10] Furthermore, according to Claim 4.2, any lower bound of the form $MA(f) \in \omega(\sqrt{n})$ either should exploit the uniformity of MA (the only difference between MA and $\widetilde{MA}$), or it will have non-trivial consequences for $AM(f)$. This partially explains why no such lower bound has been found yet.

Definition 4.1 (*Non-uniform Merlin-Arthur, $\widetilde{MA}$*)**.** For every $n \in \mathbb{N}$, let $f : \{0,1\}^n \times \{0,1\}^n \to \{0,1,\perp\}$.

If for some $k(n)$, every input distribution μ_n and every $\varepsilon > 0$ there are functions $h_1, \ldots, h_{2^{k(n)}} : \{0,1\}^n \times \{0,1\}^n \to \{0,1,\perp\}$, whose P-complexity is in $O\left(k(n) \cdot \log \frac{1}{\varepsilon}\right)$, such that

$$\Pr_{(X,Y)\sim\mu_n}\left[f(X,Y) \neq \bigvee_{i=1}^{2^{k(n)}} h_i(X,Y)\right] \leq \varepsilon,$$

then we say that the $\widetilde{MA}$-complexity of f, denoted by $\widetilde{MA}(f)$, is in $O(k(n))$.

The intuition behind the above formulation is the following.[11] Any MA-protocol allows for error reduction at the cost of (at most) a multiplicative factor of $O\left(\log \frac{1}{\varepsilon}\right)$, where ε is the "target error." It has been intuitively clear that this this property of MA is important: in particular, it was used by Klauck [13] to prove his lower bound on the MA-complexity. The concept of $\widetilde{MA}$-complexity *isolates* this property, effectively allowing for its direct analysis, which is the main subject of this part of our work.

First of all, let us see that the non-uniformity is the only possible source of advantage of $\widetilde{MA}$ over MA.

Claim 4.2. *For every $n \in \mathbb{N}$, let $g_1, \ldots, g_{2^{k(n)}} : \{0,1\}^n \times \{0,1\}^n \to \{0,1,\perp\}$ be such that for every input distribution μ_n and every $\varepsilon > 0$ the conditions of Definition 4.1 hold, as well as the additional requirement that*

$$\forall i \in \left[2^{k(n)}\right] : \Pr_{(X,Y)\sim\mu_n}[h_i(X,Y) \neq g_i(X,Y)] \leq \varepsilon.$$

Then the MA-complexity of f is in $O(k(n))$.

[10] It is not too hard to demonstrate $AM(f) \in \Omega(\log n)$ for an explicit f: for example, it holds for so-called index function $Ind(x,i) \stackrel{\text{def}}{=} x_i$; however, it is not clear how to use such examples to obtain $\widetilde{MA}(f) \in \Omega\left(\sqrt{n\log n}\right)$.

[11] The author thanks the anonymous referee whose comment has resulted in the appearance of this paragraph.

Note that the functions $g_1,\dots,g_{2^{k(n)}}$ are fixed (for every n), in particular, they do not depend on μ_n ε. The statement says that in order to become sufficient for MA, the definition of $\widetilde{MA}$ should be restricted by the additional requirement that all the h_i are approximations of the corresponding g_i. That is why we view $\widetilde{MA}$ as a non-uniform modification of MA.

Proof. Assume $\widetilde{MA}(f) \in O(k(n))$. For every input distribution μ and $\varepsilon > 0$, let $h_i^{\mu,\varepsilon}$ denote the function h_i corresponding to these μ and ε from the definition of $\widetilde{MA}(f)$.

Let ν be the uniform input distribution, then

$$\forall x,y: \bigvee_{i=1}^{2^{k(n)}} h_i^{\nu,\varepsilon}(x,y) \underset{\varepsilon\to 0}{\longrightarrow} \bigvee_{i=1}^{2^{k(n)}} g_i(x,y)$$

and

$$\forall x,y: \bigvee_{i=1}^{2^{k(n)}} h_i^{\nu,\varepsilon}(x,y) \underset{\varepsilon\to 0}{\longrightarrow} f(x,y),$$

therefore

$$f(x,y) \equiv \bigvee_{i=1}^{2^{k(n)}} g_i(x,y). \tag{4.1}$$

By the definition of $\widetilde{MA}$ it must hold that $P(h_i^{\mu,\frac{1}{3}}) \in O(k(n))$ for every input distribution μ. On the other hand,

$$\forall \mu: \underset{(X,Y)\sim\mu}{\boldsymbol{Pr}}\left[h_i^{\mu,\frac{1}{3}}(X,Y) \neq g_i(X,Y)\right] \leq \frac{1}{3},$$

which means that $BPP(g_i) \in O(k(n))$. Together with (4.1) this implies $MA(f) \in O(k(n))$. □

Next we claim that a super-$\sqrt{n}$ lower bound on $\widetilde{MA}(f)$ would have non-trivial consequences for $AM(f)$.

Theorem 4.3. *For any bipartite Boolean function f:*

$$\widetilde{MA}(f) \in O\left(\sqrt{n\cdot AM(f)}\right).$$

Proof. Let $AM(f) = k(n)$, that is, for every input distribution ν there is an NP-protocol of cost at most $k(n)$ that solves f with error at most $\frac{1}{3}$ with respect to ν. Via the standard accuracy amplification technique this implies that for any input distribution ν and $\varepsilon > 0$ there is an NP-protocol of cost $O\left(k(n)\cdot\log\frac{1}{\varepsilon}\right)$ that solves f with error at most ε with respect to ν. In particular, for every input distribution ν there is an NP-protocol Π_ν of cost $O\left(\sqrt{n\cdot k(n)}\right)$ that solves f with error at most $2^{-\sqrt{n/k(n)}}$ with respect to ν.

Let us see that

$$\widetilde{MA}(f) \in O\left(\sqrt{n\cdot k(n)}\right). \tag{4.2}$$

For $n \in \mathbb{N}$, take any input distribution μ_n and any $\varepsilon > 0$.

If $\varepsilon \leq 2^{-\sqrt{n/k(n)}}$, then let $h_1 = f$: as its P-complexity is at most $n \in O\left(\sqrt{n \cdot k(n)} \cdot \log \frac{1}{\varepsilon}\right)$, the "decomposition"

$$f(X,Y) = \bigvee_{i=1}^{1} h_i(X,Y)$$

satisfies the requirements of Definition 4.1 with respect to (4.2).

Now suppose that $\varepsilon > 2^{-\sqrt{n/k(n)}}$. Then Π_μ is an NP-protocol of cost $O\left(\sqrt{n \cdot k(n)}\right)$ that solves f with error less than ε with respect to μ. Let $K \in 2^{O\left(\sqrt{n \cdot k(n)}\right)}$ be the number of rectangles contained in Π_μ, denote their characteristic functions by $h_1, \ldots, h_K$. As the P-complexity of every such h_i is 1 and

$$\Pr_{(X,Y)\sim\mu_n}\left[f(X,Y) \neq \bigvee_{i=1}^{K} h_i(X,Y)\right] = \Pr_{\mu_n}\left[f(X,Y) \neq \Pi_\mu(X,Y)\right] < \varepsilon,$$

the requirements of Definition 4.1 with respect to (4.2) are satisfied, and the result follows. □

5 Conclusions

Among those communication complexity regimes that reside well beyond our current level of understanding, the model of *Arthur-Merlin (AM)* may be the closest to us. The motivation of this work has been to explore the "neighbourhood" of *AM* that we might be able to analyse.

- We have defined and analysed a new communication complexity class, *SLAM*, strictly included in *AM* and strictly stronger than the union of all previously known subclasses of *AM*.
- We have identified one possible source of hardness in proving $\omega(\sqrt{n})$ lower bounds against *MA*: such a bound would either be of a "very special form," or imply a non-trivial lower bound against *AM*.

A few questions that have remained open can be viewed as possible further steps towards understanding *AM*. For instance:

- What is the *SLAM*-complexity of *Disj*? Note that even its *UAM*-complexity is not understood yet (see [10] for details).
- Can we prove a lower bound of $\Omega\left(\sqrt{n \log n}\right)$ on the *MA*-complexity of an explicit function (see Footnote 10)?
- What approaches to understanding *AM* look promising?
  - Shall we try hard to prove a lower bound of $n^{1/2+\Omega(1)}$ on the *MA*-complexity of an explicit function?

– Are there complexity classes inside *AM* with non-trivial advantage over *SLAM* (or incomparable to it), which we can analyse?

Finally, we would like to mention a result that is somewhat dual to this work from the conceptual point of view: Bouland, Chen, Holden, Thaler and Vasudevan [5] define communication complexity classes $NISZK^{cc}$ and SZK^{cc}, and show that

$$NISZK^{cc} \subseteq SZK^{cc} \subseteq AM \quad \text{and} \quad NISZK^{cc} \not\subseteq UPP,$$

where *UPP* is the class consisting of functions with *high sign-rank*; *UPP* is known to strictly contain *PP*. Accordingly, the quest of understanding *AM* is at least as hard as that of understanding $NISZK^{cc}$, and the latter might be simpler if $NISZK^{cc} \subset AM$.

Acknowledgements

I am grateful to Thomas Watson for many useful comments and suggestions. I would also like to thank the authors of helpful anonymous reviews. The most careful proofreading done by the editorial team of *Theory of Computing* has contributed not just to the presentation quality of the results, but also to the author's confidence in them.

References

[1] SCOTT AARONSON AND AVI WIGDERSON: Algebrization: A new barrier in complexity theory. *ACM Trans. Comput. Theory*, 1(1):1–54, 2009. Preliminary version in STOC'08. [doi:10.1145/1490270.1490272] 4, 23

[2] HAROLD ABELSON: Lower bounds on information transfer in distributed computations. *J. ACM*, 27(2):384–392, 1980. Preliminary version in FOCS'78. [doi:10.1145/322186.322200] 6

[3] LÁSZLÓ BABAI, PÉTER FRANKL, AND JANOS SIMON: Complexity classes in communication complexity theory. In *Proc. 27th FOCS*, pp. 337–347. IEEE Comp. Soc., 1986. [doi:10.1109/SFCS.1986.15] 6

[4] ELMAR BÖHLER, CHRISTIAN GLASSER, AND DANIEL MEISTER: Error-bounded probabilistic computations between MA and AM. *J. Comput. System Sci.*, 72(6):1043–1076, 2006. Preliminary version in MFCS'03. [doi:10.1016/j.jcss.2006.05.001] 7

[5] ADAM BOULAND, LIJIE CHEN, DHIRAJ HOLDEN, JUSTIN THALER, AND PRASHANT VASUDEVAN: On the power of statistical zero knowledge. *SIAM J. Comput.*, 49(4):1–58, 2020. Preliminary version in FOCS'17. [doi:10.1137/17M1161749, arXiv:1609.02888] 26

[6] LIJIE CHEN: On the hardness of approximate and exact (bichromatic) maximum inner product. *Theory of Computing*, 16(4):1–50, 2020. Preliminary version in CCC'18. [doi:10.4086/toc.2020.v016a004] 4, 23

[7] EVGENY DRUKH AND YISHAY MANSOUR: Concentration bounds for unigram language models. *JMLR*, 6(42):1231–1264, 2005. JMLR. 20

[8] SHAFI GOLDWASSER AND MICHAEL SIPSER: Private coins versus public coins in interactive proof systems. In SILVIO MICALI, editor, *Randomness in Computation*, volume 5 of *Advances in Computing Research*, pp. 73–90. JAI Press, 1989. Preliminary version in STOC'86. 7, 20

[9] MIKA GÖÖS, SHACHAR LOVETT, RAGHU MEKA, THOMAS WATSON, AND DAVID ZUCKERMAN: Rectangles are nonnegative juntas. *SIAM J. Comput.*, 45(5):1835–1869, 2016. Preliminary version in STOC'15. [doi:10.1137/15M103145X] 3, 7, 8, 21, 22

[10] MIKA GÖÖS, TONIANN PITASSI, AND THOMAS WATSON: Zero-information protocols and unambiguity in Arthur–Merlin communication. *Algorithmica*, 76(3):684–719, 2016. Preliminary version in ITCS'15. [doi:10.1007/s00453-015-0104-9] 3, 4, 8, 10, 21, 22, 25

[11] MIKA GÖÖS, TONIANN PITASSI, AND THOMAS WATSON: The landscape of communication complexity classes. *Comput. Complexity*, 27(2):245–304, 2018. [doi:10.1007/s00037-018-0166-6] 3, 5

[12] MAURICIO KARCHMER, ILAN NEWMAN, MIKE SAKS, AND AVI WIGDERSON: Non-deterministic communication complexity with few witnesses. *J. Comput. System Sci.*, 49(2):247–257, 1994. Preliminary version in SCT(CCC)'92. [doi:10.1016/S0022-0000(05)80049-2] 8

[13] HARTMUT KLAUCK: Rectangle size bounds and threshold covers in communication complexity. In *Proc. 18th IEEE Conf. on Comput. Complexity (CCC'03)*, pp. 118–134. IEEE Comp. Soc., 2003. [doi:10.1109/CCC.2003.1214415, arXiv:cs/0208006] 2, 4, 10, 23

[14] HARTMUT KLAUCK: On Arthur Merlin games in communication complexity. In *Proc. 26th IEEE Conf. on Comput. Complexity (CCC'11)*, pp. 189–199. IEEE Comp. Soc., 2011. [doi:10.1109/CCC.2011.33, arXiv:1101.0523] 3, 7, 12, 20

[15] EYAL KUSHILEVITZ AND NOAM NISAN: *Communication Complexity*. Cambridge Univ. Press, 1997. [doi:10.1017/CBO9780511574948] 5

[16] JOHN VON NEUMANN: Zur Theorie der Gesellschaftsspiele. *Mathematische Annalen*, 100(1):295–320, 1928. EuDML. 6, 20

[17] ANDREW CHI-CHIH YAO: Some complexity questions related to distributive computing. In *Proc. 11th STOC*, pp. 209–213. ACM Press, 1979. [doi:10.1145/800135.804414] 6

AUTHOR

Dmitry Gavinsky
Researcher
Department of Mathematical Logic
and Theoretical Computer Science
Institute of Mathematics of the
Czech Academy of Sciences
Praha, Czech Republic
gavinsky@math.cas.cz
https://users.math.cas.cz/~gavinsky/

ABOUT THE AUTHOR

In the good old days DMITRY GAVINSKY studied at the Technion - Israel Institute of Technology and at the University of Calgary. Thanks to the support of his scientific advisers Nader Bshouty, Richard Cleve and John Watrous, he graduated from both institutions, and now he lives and works in the best city in the world. There he enjoys the best beer and loves *good* music and *good* books.