Verifiable Stream Computation and Arthur-Merlin Communication

Amit Chakrabarti[1], Graham Cormode[2], Andrew McGregor[3], Justin Thaler[4], and Suresh Venkatasubramanian[5]

1   Department of Computer Science, Dartmouth College, USA
2   Department of Computer Science, University of Warwick, UK
3   Department of Computer Science, UMass Amherst, USA
4   Yahoo Labs, New York, USA
5   School of Computing, University of Utah, USA

Abstract

In the setting of streaming interactive proofs (SIPs), a client (verifier) needs to compute a given function on a massive stream of data, arriving online, but is unable to store even a small fraction of the data. It outsources the processing to a third party service (prover), but is unwilling to blindly trust answers returned by this service. Thus, the service cannot simply supply the desired answer; it must *convince* the verifier of its correctness via a short interaction after the stream has been seen.

In this work we study "barely interactive" SIPs. Specifically, we show that two or three rounds of interaction suffice to solve several query problems – including Index, Median, Nearest Neighbor Search, Pattern Matching, and Range Counting – with polylogarithmic space and communication costs. Such efficiency with $O(1)$ rounds of interaction was thought to be impossible based on previous work.

On the other hand, we initiate a formal study of the limitations of constant-round SIPs by introducing a new hierarchy of communication models called Online Interactive Proofs (OIPs). The online nature of these models is analogous to the streaming restriction placed upon the verifier in an SIP. We give upper and lower bounds that (1) characterize, up to quadratic blowups, every finite level of the OIP hierarchy in terms of other well-known communication complexity classes, (2) separate the first four levels of the hierarchy, and (3) reveal that the hierarchy collapses to the fourth level. Our study of OIPs reveals marked contrasts and some parallels with the classic Turing Machine theory of interactive proofs, establishes limits on the power of existing techniques for developing constant-round SIPs, and provides a new characterization of (non-online) Arthur–Merlin communication in terms of an online model.

1   Introduction

The surging popularity of commercial cloud computing services, and more generally outsourced computations, has revealed compelling new applications for the study of *interactive proofs* with highly restricted *verifiers*. Consider, e.g., a retailer (verifier) who lacks the resources to locally process a massive input (say, the set of all its transactions), but can access a powerful but untrusted cloud service provider (prover), who processes the input on the retailer's behalf. The verifier must work within the confines of the restrictive *data streaming* paradigm, using

only a small amount of working memory. The prover must both answer queries about the input (say, "how many pairs of blue jeans have I ever sold?"), and prove that the answer is correct. We refer to this general scenario as *verifiable data stream computation*.

It is useful to look at this computational scenario as "data stream algorithms with access to a powerful (space-unlimited) prover." As is well known, most interesting data streaming problems have no nontrivial (i.e., sublinear space) algorithms unless one allows approximation. For instance, given a stream σ of tokens from the universe $[n] := \{1, 2, \ldots, n\}$, which implicitly defines frequencies f_j for each $j \in [n]$, some basic questions we can ask about σ are the number of distinct tokens $F_0(\sigma)$, the kth frequency moment $F_k(\sigma) = \sum_{j=1}^n f_j^k$, the median of the collection of numbers in σ, and the very basic *point queries* where, given a specific $j \in [n]$ after σ has been presented, we wish to know f_j. In each case, we would like an *exact* answer, not an estimate. With the trivial exception of $F_1(\sigma)$ – which is just the length of σ – not one of these questions can be answered by a (possibly randomized) streaming algorithm restricted to $o(n)$ space. However, with access to a powerful prover, things improve greatly: as shown in Chakrabarti et al. [9], point queries, median, and F_k (for integral $k > 0$) can be computed exactly by a verifier using $\tilde{O}(\sqrt{n})$ space, while receiving $\tilde{O}(\sqrt{n})$ bits of "help" from the prover.

Notably, the protocol achieving this $\tilde{O}(\sqrt{n})$ cost (space plus amount of help) is non-interactive: the prover sends a single message to the verifier. Chakrabarti et al. [9] also showed that under this restriction their protocol is optimal: a cost of $\Omega(\sqrt{n})$ is required. In subsequent work, Cormode et al. [15] considered *streaming interactive proofs* (SIPs), where the verifier may engage in several rounds of interaction with the prover, seeking to minimize both the space used by the verifier and the total amount of communication. They gave SIPs with $2k - 1$ rounds of interaction following the verifier's single pass over the input stream, achieving a cost of $\tilde{O}(n^{1/(k+1)})$ for the above problems. This generalizes the earlier set of results [9], which gave 1-round SIPs. Moreover, it achieves $O(\text{polylog}\, n)$ cost with $O(\log n / \log \log n)$ rounds of interaction. In recent work, Klauck and Prakash [24] further studied this kind of computation and generalized the 1-round lower bound, claiming that a $(2k - 1)$-round SIP must cost $\Omega(n^{1/(k+1)})$, even for very basic point queries.

However, we identify an implicit assumption in the Klauck–Prakash lower bound argument: it applies only to protocols in which the verifier's messages to the prover are independent of the input. This happened to hold in all previous SIPs, which are ultimately descended from the sum-check protocol of Lund et al. [27]. Furthermore, this assumption is harmless in the classical theory of interactive proofs where public-coin protocols can simulate private-coin ones with just a polynomial blowup in cost [16]. However, these simulation results fail subtly in the streaming setting, and we show that this failure is intrinsic by giving a number of new upper bounds.

1.1   New Results: Exponentially Improved Constant-Round SIPs

We start by showing that even two-round SIPs are exponentially more powerful than previously believed, on certain problems. For now we state our results informally, using the $\tilde{O}$-notation to suppress "lower order" factors. We give formal theorem statements later in the paper, after all definitions are in place.

▶ **Result 1.1** (Formalized in Theorem 3.1). *There is a two-round SIP with cost $\tilde{O}(\log n)$ for answering point queries on a stream over the universe $[n]$.*

The SIP that achieves this upper bound is based on an abstract protocol that we call the *polynomial evaluation protocol.* Crucially, unlike the sum-check protocols used in previous

SIPs, it involves an interaction where the verifier's message to the prover depends on part of the input; specifically, it depends on the query. Note that two rounds of interaction is likely unavoidable in practice even if verifiability is not a concern: one round may be required for the verifier to communicate the query to the prover, with a second round required for the prover to reply.

Adding a third round of interaction allows us to answer selection queries, of which an important special case is median-finding.

▶ **Result 1.2** (Formalized in Theorem 3.7). *There is a three-round SIP with cost $\tilde{O}(\log n)$ for determining the exact median of a stream of numbers from $[n]$.*

We can in fact answer fairly complex queries with three rounds and polylogarithmic cost. For instance, given a data set presented as a stream of points from a metric space, we can answer *exact* nearest neighbor queries to the data set very efficiently, even in high dimensions. This is somewhat surprising, given that even the offline version of the problem seems to exhibit a curse of dimensionality.

▶ **Result 1.3** (Formalized in Theorem 3.4). *For data sets consisting of points from $[n]^d$ under a reasonable metric, such as the Manhattan distance ℓ_1^d or the Euclidean distance ℓ_2^d, there is a three-round SIP with cost $\mathrm{poly}(d, \log n)$ allowing exact nearest neighbor queries to the data set.*

We also give similarly efficient two-round SIPs for other well-studied query problems, such as range counting queries (Theorem 3.6), where a stream of data points is followed by a query range and the goal is to determine the number of points in the range that appeared in the stream, and pattern matching queries (Theorem 3.8), where a streamed text is followed by a (short) query pattern.

Next, we work towards a detailed understanding of the subtleties of SIPs that caused the aforementioned Klauck–Prakash lower bound [24] not to apply. Our study naturally leads into communication complexity, in particular to Arthur–Merlin communication, which we discuss next.

1.2 The Connection to Arthur–Merlin Communication

Like almost all previous lower bounds for data stream computations, prior SIP lower bounds [9, 24] use reductions from problems in communication complexity. To model the prover in an SIP, the appropriate setting is Arthur–Merlin communication, which we now introduce.

Suppose Alice holds an input $x \in \mathcal{X}$, Bob holds $y \in \mathcal{Y}$, and they wish to compute $f(x, y)$ for some Boolean function $f : \mathcal{X} \times \mathcal{Y} \to \{0, 1\}$, using random coins and settling for some constant probability of error. Say this costs $\mathrm{R}(f)$ bits of communication. Can an omniscient but untrusted Merlin, who knows (x, y), convince "Arthur" (defined as Alice and Bob together) that $f(x, y) = 1$, keeping the overall communication within $o(\mathrm{R}(f))$? For several interesting functions f the answer is "Yes" and this is the general subject of Arthur–Merlin communication complexity, first considered in seminal work by Babai, Frankl, and Simon [5].

The one-pass streaming restriction on the verifier in an SIP is modeled by requiring that Alice not receive any communication from either Bob or Merlin. Thus the Alice–Bob communication is one-way, though Bob and Merlin may interact arbitrarily. We refer to this restricted communication setting as *online Arthur–Merlin communication*. It should be clear that a k-round SIP with cost C can be simulated by an online Arthur–Merlin communication

of cost C where Bob and Merlin interact for k rounds. Thus, lower bounds on SIPs would follow from corresponding communication lower bounds in the online Arthur–Merlin setting.

At this point let us recall that the classical Turing-Machine-based theory of interactive proofs considers two different models of interaction between prover and verifier, corresponding to the complexity classes $\mathbf{IP_{TM}}$,[1] where the verifier is allowed private randomness, and $\mathbf{AM_{TM}}$, where he may only use public randomness. Recall the following classic results about such interactive proofs.

- **Equivalence of private and public coins.** Goldwasser and Sipser [16] proved that a k-round private coin interactive proof (à la $\mathbf{IP_{TM}}$) can be simulated (with a polynomial blowup in complexity) by a $(k + 2)$-round public coin one (à la $\mathbf{AM_{TM}}$). Thus, in the resulting protocol, the verifier can perform his interaction with the prover before even looking at the input!

- **Round reduction.** Babai and Moran [6] proved that a $(k + 1)$-round interactive proof can be simulated by a k-round interactive proof with a polynomial blowup in the verifier's complexity. Thus, a two-round (verifier→prover→verifier) interactive proof is just as powerful as any constant-round one.

Interestingly, as we shall show in this work, neither of these phenomena holds for the *online* communication complexity analogs of $\mathbf{IP_{TM}}$ and $\mathbf{AM_{TM}}$. (Recall that "online" means that Alice does not receive any communication from either Bob or Merlin.) This point appears to have been missed in the Klauck–Prakash proof [24], which works in a "public coin" setting and thus applies only to a restricted class of SIPs. The new SIPs we design in this work correspond to a "private coin" setting, which allows the aforementioned exponential improvements.

Clearly there are nuances in online Arthur–Merlin communication complexity that do not arise in classical interactive proofs. In particular, we seek a better understanding of the precise role of rounds and of private randomness in the communication setting. This is the goal of our next batch of results.

1.3 New Results: Complexity Classes for Arthur–Merlin Communication

As noted above, we can think of $\mathbf{AM_{TM}}$ as a restricted interactive proof model where the verifier must interact with the prover before looking at his input. We can then define a hierarchy of analogous communication complexity models called $\mathbf{OMA^{[k]}}$ (Online Merlin–Arthur), where Bob interacts with Merlin in k rounds without looking at his input, and then Alice communicates with Bob one-way. We defer precise definitions to Section 4. The aforementioned Klauck–Prakash lower bound essentially says the following:

▶ **Proposition 1.4** (Klauck and Prakash [24]). *The* INDEX *problem – where Alice gets $x \in \{0,1\}^n$, Bob gets $j \in [n]$ and Bob must output x_j with high probability – requires $\Omega(n^{1/(k+1)})$ cost in the* $\mathbf{OMA^{[2k]}}$ *model.*

We can also define a parallel hierarchy $\mathbf{OIP^{[k]}}$ (Online Interactive Proof) of communication analogs of $\mathbf{IP_{TM}}$. We now hit another subtlety. We could require the Bob–Merlin interaction to happen before the Alice→Bob communication; this is how we shall define $\mathbf{OIP^{[k]}}$. Alternatively, we could swap the order, so that Bob's messages to Merlin could depend on Alice's input as well; we shall call the resulting (more powerful) model $\mathbf{OIP_+^{[k]}}$.

---

[1] Throughout this paper, we use the subscript "TM" to denote a Turing-machine-based complexity class, to resolve the notation clash with the analogous communication complexity classes.

These communication models correspond to SIPs as follows. Every SIP designed prior to this work falls into a restricted setting where the verifier's messages are independent of the input, so it can be simulated by an $\mathbf{OMA^{[k]}}$ protocol with k being the number of rounds of interaction in the SIP. The SIPs we design in this work apply to "query problems" with the data set appearing before the query, and our verifier messages depend only on the query. Thus our SIPs are naturally simulable by $\mathbf{OIP^{[k]}}$ protocols. Finally, a general SIP, where verifier messages can depend on the entire input stream, is simulable by an $\mathbf{OIP_+^{[k]}}$ protocol.

Following Babai et al. [5], given a communication model $\mathbf{C}$, we define a corresponding complexity class, also denoted $\mathbf{C}$, consisting of all problems that have *polylogarithmic* cost protocols in the model $\mathbf{C}$. We now have three parallel hierarchies of communication complexity classes: $\mathbf{OMA^{[k]}}$, $\mathbf{OIP^{[k]}}$, and $\mathbf{OIP_+^{[k]}}$. For our next batch of results, we prove several inclusion and separation results relating these newly defined classes to each other and to well-studied classes from earlier work in communication complexity.

▶ **Result 1.5** (Formalized over several theorems in Section 5). *The following complexity class inclusions and separations, given in Figure 1, hold.*

■ **Figure 1** The layout of our communication complexity zoo. An arrow from $\mathbf{C}_1$ to $\mathbf{C}_2$ indicates that $\mathbf{C}_1 \subseteq \mathbf{C}_2$. If the arrow is double-headed, then the inclusion is strict. Here $k > 4$ is an arbitrary constant. The models $\mathbf{R^{[t,A]}}$ (resp. $\mathbf{R^{[t,B]}}$) are standard t-round randomized communication with Alice (resp. Bob) starting. The model $\mathbf{MA^{[2,B]}}$ consists of a message from Merlin followed by Bob→Alice→Bob communication, while $\mathbf{AM}$ is standard (see Section 5).

Notice that there are several two-way inclusions (i.e., equalities) amongst these communication complexity classes. It is worth noting that with one exception (namely $\mathbf{OIP^{[1]}} = \mathbf{OIP_+^{[1]}}$) none of these equalities is trivial. For instance, consider the switch from the model $\mathbf{R^{[2,B]}}$ to the model $\mathbf{OIP^{[2]}}$: Bob loses the ability to send Alice a message before hearing from her, but gains access to Merlin. It is not *a priori* clear that this switch in models will result in a complexity class that is even comparable to $\mathbf{R^{[2,B]}}$, and nontrivial simulation arguments (Theorems 5.3 and 5.6) are required to prove that $\mathbf{R^{[2,B]}} = \mathbf{OIP^{[2]}}$.

Many of our simulations incur some blowup in cost. All such blowups are at most quadratic, so polylogarithmic costs remain polylogarithmic.

The $\mathbf{OMA}$ and $\mathbf{OIP}$ hierarchies behave quite differently from the classical $\mathbf{AM_{TM}}$ and $\mathbf{IP_{TM}}$:

- In contrast to the Goldwasser–Siper private-by-public-coin theorem, the class $\mathbf{OIP^{[4]}}$ is strictly more powerful than $\mathbf{OMA^{[k]}}$ (in fact, even $\mathbf{OIP^{[2]}} \not\subseteq \mathbf{OMA^{[k]}}$), for every constant k.
- In contrast to the Babai–Moran round reduction theorem, there are exactly four distinct levels (not two) in the $\mathbf{OIP^{[k]}}$ hierarchy, for constant k.

In the course of proving the separation results in Figure 1, we obtain concrete lower bounds for explicit functions that are of interest in their own right. Let us highlight one of these.

▶ **Result 1.6** (Formalized in Corollary 5.9). *The set disjointness problem* DISJ *– where Alice and Bob each get a subset of* $[n]$ *and must decide whether they are disjoint – requires* $\Omega(n^{1/3})$ *cost in the* $\mathbf{OIP}^{[3]}$ *model and thus does not belong to the class* $\mathbf{OIP}^{[3]}$. *This lower bound is tight up to a logarithmic factor.*

This has implications for SIPs. We noted that all SIPs designed thus far (including the new ones in this work) are simulable in the weaker **OIP** models. By a standard reduction [4] from DISJ to the frequency moments problem F_k, it follows that unlike what we achieved for point queries and median queries, *based on currently known techniques*, we cannot get a polylogarithmic cost three-round SIP for F_k ($k \neq 1$).

Removing the qualifier "based on currently known techniques" above would require a similar lower bound for $\mathbf{OIP}^{[3]}_+$. Unfortunately, at present we are unable to prove any nontrivial lower bounds on $\mathbf{OIP}^{[2]}_+$, and doing so appears to be a rather difficult problem. Indeed, this inability is what led us to study the weaker **OIP** model. Yet, because the **OIP** models are online, the separation results in Figure 1 still morally capture the primary way in which SIPs, due to their streaming/online nature, differ from classical interactive proofs.

Finally, our result $\mathbf{AM} = \mathbf{OIP}^{[4]}$ gives a novel characterization of **AM** in terms of *online* communication. This is surprising because online models, where no one talks to Alice, might be expected to be too weak to capture **AM**. Proving lower bounds on **AM** is a longstanding and notoriously hard problem in communication complexity [26, 22, 23]. We believe our new characterization of **AM** is of independent interest, and may prove useful in establishing non-trivial **AM** lower bounds.

1.4 Related Work

1.4.1 Stream Computation

Early theoretical work on verifiable stream computation focused on non-interactive protocols, as in the *annotated data streams* model of Chakrabarti et al. [9]. In our language, that model corresponds to 1-round SIPs. Work in this model has established optimal protocols for several problems including frequency moments and frequent items [9]; linear algebraic problems such as matrix rank [24]; and graph problems like shortest s–t path [14]. Many of these protocols have subsequently been optimized for streams whose length is much smaller than the universe size [8]. More recent protocols, such as the *Arthur–Merlin streaming protocols* of Gur and Raz [19, 8] are "barely interactive" in the sense that the prover and the verifier may exchange a constant number of messages. Meanwhile, the fully general *streaming interactive proof* (SIP) model of Cormode et al. [15, 13] permits "many" rounds of interaction. Cormode, Thaler, and Yi [15] showed that several general $\mathbf{IP}_{\mathsf{TM}}$ protocols can be simulated in this model. These include the powerful, general-purpose protocol of Goldwasser, Kalai, and Rothblum [18]. Given any problem in $\mathbf{NC}_{\mathsf{TM}}$, the resulting protocol requires only polylogarithmic space and communication while using polylogarithmic rounds of verifier–prover interaction. Refinements and implementations of these protocols [36, 13, 35] have demonstrated scalability and the practicality of this line of work.

Algebraic techniques lie at the core of almost all nontrivial protocols in the above models. Specifically, a number of 1-round SIPs are inspired by the Aaronson–Wigderson **MA** communication protocol for DISJ [2], which is in turn inspired by the classic sum-check

protocol of Lund et al. [27]. The sum-check protocol is also the inspiration for the way that all previous multi-round SIPs make use of interaction. The aforementioned protocol of Goldwasser et al. [18] also builds upon the sum-check protocol.

The algorithmic results outlined in Section 1.1 have a rather different algebraic idea at their core. They are based on the aforementioned *polynomial evaluation protocol*, which is obtained by adapting a result of Raz [33] about $\mathbf{IP/rpoly}_{\mathsf{TM}}$ to a streaming setting; see the discussion at the start of Section 2.1.

Early work on interactive proofs studied space-bounded verifiers (see the survey by Condon [12]), but many protocols developed in this line of work require the verifier to store the input, and therefore do not address verifiable *stream* computation, as we do here. Goldwasser et al. [17] studied interactive proofs with verifiers in the complexity class $\mathbf{NC}^0_{\mathsf{TM}}$. Interestingly, they showed that private randomness is necessary to obtain interactive proofs with verifiers in $\mathbf{NC}^0_{\mathsf{TM}}$, unless the language in question is already in $\mathbf{NC}^0_{\mathsf{TM}}$. This is analogous to our finding that constant-round "public coin" SIPs (where the verifier's messages do not depend on the input) are exponentially weaker than general constant-round SIPs.

1.4.2   Computationally Sound Protocols

Protocols for verifiable stream computation have also been studied in the cryptography community [10, 32, 34]. These works only require soundness to hold against cheating provers that run in polynomial time. In exchange for this weaker security guarantee, these protocols can achieve properties that are impossible in the information-theoretic setting we consider. For example, they typically achieve *reusability*, allowing the verifier to use the same randomness to answer many queries. In contrast, our protocols only support "one-shot" queries, because they require the verifier to reveal secret randomness to the prover.

Chung et al. [10] combine the GKR protocol with fully homomorphic encryption (FHE) to give reusable, non-interactive protocols of polylogarithmic cost for any problem in $\mathbf{NC}$. Papamanthou et al. [32] give improved protocols for a class of low-complexity queries including point queries and range search: their protocols avoid the use of FHE, and allow the prover to answer such queries in polylogarithmic time (a similar property was achieved by Schröder and Schröder [34], but for a simpler class of queries, and with unidrectional communication from the verifier to the prover on each stream update). In contrast, prior work as well as our own requires the prover to spend time quasilinear in the size of the data stream after receiving a query, even if the answer itself can be computed in sublinear time (e.g., point queries, which can be solved with a single access to memory). We note however that our most interesting protocols, such as those for nearest neighbor search and pattern matching, are for problems that cannot be solved in sublinear time; hence, the quasilinear time required by our protocols does not affect the prover's runtime by more than logarithmic factors.

1.4.3   Communication Complexity

Seminal work by Babai et al. [5] introduced and studied the communication analogs of the major Turing Machine complexity classes, including $\mathbf{P}$, $\mathbf{NP}$, $\mathbf{\Sigma}_2$, $\mathbf{\Pi}_2$. They hinted at similar analogs of $\mathbf{MA}$ and the $\mathbf{AM}$ hierarchy. Lokam [26] related the task of placing problems outside of the communication class $\mathbf{AM}$ to notions of matrix rigidity. He also observed that the communication complexity classes $\mathbf{IP}$ and $\mathbf{AM}$ behave similarly to their Turing Machine counterparts. In particular, noted theorems such as $\mathbf{IP} = \mathbf{PSPACE}$, Toda's Theorem, and

Babai and Moran's round reduction results [6] all hold in the communication world (though not under *online* communication, as shown by this work).

Online (also known as one-round) randomized communication complexity was introduced in the mid-1990s and considered by Ablayev [3], Kremer, Nisan, and Ron [25], and Newman and Szegedy [29]. Aaronson [1] introduced online variants of Merlin–Arthur communication, in classical and quantum flavors. Aaronson and Wigderson [2] gave an online **MA** communication protocol for DISJ (more generally, for INNER-PRODUCT) with cost $\tilde{O}(\sqrt{n})$; this is nearly optimal, as shown by a lower bound of Klauck [22] that applies to general **MA** protocols. More recently, Klauck [23] performed a careful study of **AM**, **MA**, and its quantum analogue **QMA**. In particular, he gave a promise problem PAPPMP separating **QMA** from **AM**; we shall eventually show that PAPPMP separates **OIP**[3] from **OIP**[4].

<div style="border-left: 4px solid orange;">

2   The SIP Model and the Polynomial Evaluation Protocol

</div>

In a data stream problem, the input σ is a *stream*, or sequence, of *tokens* from some data universe $\mathcal{U}$. The goal is to compute or approximate some function $g(\sigma)$, keeping space usage sublinear in the two key size parameters: (1) the length of σ, and (2) the size of the universe $|\mathcal{U}|$. Practically speaking, we would also like to process each stream update (token arrival) quickly. All our data stream algorithms will be randomized, and we shall allow them to err with some small constant probability on each input stream. In the *streaming interactive proofs* (SIP) model, after processing σ, the algorithm (called the "verifier") may engage in k rounds of interaction with an oracle (the "prover") who knows σ and whose goal is to lead the verifier to output the correct answer $g(\sigma)$. The verifier, being distrustful, will output "$\perp$" (indicating "abort") if he suspects the prover to be cheating.

All of the SIPs in this paper will work in the *turnstile streaming model*, where σ can contain deletions of tokens from $\mathcal{U}$, in addition to insertions. In this model it is best to think of the input as being a stream of integer *updates* to a vector $\mathbf{x} = (x_1, \ldots, x_n) \in \mathbb{Z}^n$. Initially $\mathbf{x} = \mathbf{0}$, and an update is a tuple $(i, c) \in [n] \times \mathbb{Z}$, which has the effect of adding c to the entry x_i. We will sometimes describe our algorithms as they apply to the vanilla streaming model, but it will be straightforward to extend them to the turnstile model.

We say that an SIP computes the function g with *completeness error* ε_c and *soundness error* ε_s if for all inputs $\mathbf{x}$ there exists a prover strategy that will cause the verifier to output $g(\mathbf{x})$ with probability at least $1 - \varepsilon_c$, and no prover strategy can cause the verifier to output a value outside $\{g(x), \perp\}$ with probability larger than ε_s. In designing SIPs, our goal will be to achieve $\varepsilon_c, \varepsilon_s \leq 1/3$; clearly the theory remains unchanged if we replace $1/3$ by another constant in $(0, 1/2)$. A SIP with $\varepsilon_c = 0$ is said to have *perfect completeness*. The total length of the verifier–prover interaction is the *help cost*. The space used by the streaming verifier is the *space cost*. The *cost* of an SIP is the sum of its help cost and its space cost. When designing SIP protocols we will also discuss the time complexities of the prover and the verifier. To keep things simple, we consider a model in which all arithmetic operations on a finite field of size $n^{O(1)}$ can be executed in unit time.

2.1   The Polynomial Evaluation Protocol

We shall present a two-round SIP for an abstract data stream problem called "polynomial evaluation," where the input consists of a multivariate polynomial described implicitly, as a table of values, followed by a point at which the polynomial must be evaluated. Without space constraints, this problem simply amounts to interpolation followed by direct evaluation, but our goal is to obtain a protocol where the verifier uses space roughly logarithmic in the

size of the table of values, and is convinced by the prover about the correct answer after a similar amount of communication. For ease of presentation, we shall first consider a concrete special setting that is important in its own right: the INDEX problem. In this problem, the input is a stream of n *data bits* $x_1, \ldots, x_n$, followed by a *query index* $j \in [n]$. The goal is to output x_j with error at most $1/3$.

With very different motivations from ours, Raz [33] gave an interactive proof protocol placing *every* language in $\mathbf{IP}_{\mathsf{TM}}/\mathbf{rpoly}$, the class of languages that have interactive proofs with polynomial-time verifiers that take randomized advice, where the advice is kept secret from the prover. Our SIP for INDEX can be seen as an adaptation of Raz's interactive proof to the streaming setting.

▶ **Theorem 2.1.** *The* INDEX *problem has a two-round SIP with cost* $O(\log n \log \log n)$, *in which the verifier processes each stream token in* $O(\log n)$ *time and the prover runs in total time* $O(n \log n)$.

Proof. Assume WLOG that $n = 2^b$, for some integer b. Identify each integer $z \in [n]$ with a Boolean vector $\mathbf{z} = (z_1, \ldots, z_b) \in \{0,1\}^b$ in some canonical way, such as by using the binary representation of z. We can then view the data bits as a table of values for the Boolean function $g_x : \{0,1\}^b \to \{0,1\}$ given by $g_x(\mathbf{z}) = x_z$, and thus for the multilinear b-variate polynomial $\widetilde{g}_x(Z_1, \ldots, Z_b)$ given by

$$\widetilde{g}_x(Z_1, \ldots, Z_b) = \sum_{\mathbf{z} \in \{0,1\}^b} g_x(\mathbf{z}) \chi_{\mathbf{z}}(Z_1, \ldots, Z_b), \text{ where} \tag{1}$$

$$\chi_{\mathbf{u}}(Z_1, \ldots, Z_b) = \prod_{i=1}^{b} \big((1 - u_i)(1 - Z_i) + u_i Z_i\big) \tag{2}$$

is the indicator function of the vector $\mathbf{u} = (u_1, \ldots, u_b)$. We shall interpret $\widetilde{g}_x$ as a polynomial in $\mathbb{F}[Z_1, \ldots, Z_b]$ for a fixed "large enough" finite field $\mathbb{F}$. With this interpretation, $\widetilde{g}_x$ is called the multilinear extension of g_x to $\mathbb{F}$. We define a *line* in $\mathbb{F}^b$ to be the range of a nonconstant affine function from $\mathbb{F}$ to $\mathbb{F}^b$. Every line contains exactly $|\mathbb{F}|$ points. Given such a line, ℓ, we define its *canonical representation* to be the degree-1 polynomial $\lambda_\ell(W) \in \mathbb{F}^b[W]$ such that $\lambda_\ell(0)$ and $\lambda_\ell(1)$ are, respectively, the lexicographically first and second points in ℓ. We define the *canonical restriction* of a polynomial $f(Z_1, \ldots, Z_b)$ to ℓ to be the univariate polynomial $f(\lambda_\ell(W)) \in \mathbb{F}[W]$, whose degree is at most the total degree of f.

Using the above notations and conventions, our two-round SIP for INDEX works as shown in Figure 2.

To analyze this protocol, first note that after reading all the data bits, the verifier would have computed $Q = \widetilde{g}_x(\mathbf{r})$, by Eq. (1). Now the protocol is easily seen to have perfect completeness. Since $\widetilde{g}_x(Z_1, \ldots, Z_b)$ is multilinear, it follows that $\deg(\widetilde{g}_x(\lambda_\ell(W))) \le b$, so the prover can always honestly choose $h(W) = \widetilde{g}_x(\lambda_\ell(W))$. If he does so, then we will indeed have $h(t) = \widetilde{g}_x(\lambda_\ell(t)) = \widetilde{g}_x(\mathbf{r}) = Q$, and the verifier's check will pass. Finally, the verifier will output $h(w) = \widetilde{g}_x(\lambda_\ell(w)) = \widetilde{g}_x(\mathbf{j}) = x_j$, the correct answer to the INDEX instance.

Next, we analyze soundness. If the prover supplies a polynomial $h(W) \not\equiv \widetilde{g}_x(\lambda_\ell(W))$, then, since both polynomials have degree at most b, they agree at at most b points in $\mathbb{F}$. From the prover's perspective after he receives the verifier's message, $\mathbf{r}$ is uniformly distributed in $\ell \setminus \{\mathbf{j}\}$. Thus, $\Pr_{\mathbf{r}}[h(t) = Q] \le b/(|\mathbb{F}| - 1) \le 1/3$.

Now we consider this protocol's costs. The verifier maintains the random point $\mathbf{r} \in \mathbb{F}^b$ and the running sum $Q \in \mathbb{F}$, using $O(b \log |\mathbb{F}|)$ space. He sends the prover ℓ, which is specified by two elements of $\mathbb{F}^b$, and receives a degree-b polynomial in $\mathbb{F}[W]$; both communications use at most $O(b \log |\mathbb{F}|)$ bits. Recalling that $|\mathbb{F}| \le 6b + 2$, we see that both space and communication costs are in $O(b \log b) = O(\log n \log \log n)$.

Input: Stream of data bits $(x_1, \ldots, x_n)$ where $n = 2^b$, followed by index $j \in [n]$.
Goal: Prover to convince Verifier to output the correct value of x_j.
Shared Agreement: Finite field $\mathbb{F}$ with $3b + 1 \leq |\mathbb{F}| \leq 6b + 2$; bijective map $u \in [n] \longleftrightarrow$ $\mathbf{u} \in \{0, 1\}^b$.

Initialization: Verifier picks $\mathbf{r} \in_R \mathbb{F}^b$ uniformly at random, sets $Q \leftarrow 0$.
Stream Processing: Upon reading x_z, where $z \in [n]$, Verifier updates $Q \leftarrow Q + x_z \chi_{\mathbf{z}}(\mathbf{r})$.
Query Handling: Upon reading the index j, Verifier interacts with Prover as follows:

1. If $\mathbf{j} = \mathbf{r}$, Verifier outputs Q as the answer. Otherwise, he sends Prover ℓ, the unique line in $\mathbb{F}^b$ through $\mathbf{j}$ and $\mathbf{r}$.
2. Prover sends Verifier a polynomial $h(W) \in \mathbb{F}[W]$ of degree at most b, claiming that it is the canonical restriction of the multilinear polynomial $\widetilde{g}_x(Z_1, \ldots, Z_b)$ to the line ℓ. That is, Prover claims that $h(W) \equiv \widetilde{g}_x(\lambda_\ell(W))$.
3. Let $w, t \in \mathbb{F}$ be such that $\lambda_\ell(w) = \mathbf{j}$ and $\lambda_\ell(t) = \mathbf{r}$. Verifier checks that $h(t) = Q$, aborting if not. If the check passes, Verifier outputs $h(w)$ as the answer.

■ **Figure 2** A Two-Round Streaming Interactive Proof (SIP) Protocol for the INDEX Problem

Finally, we consider the verifier's and prover's runtimes. The honest prover must send the univariate polynomial $\widetilde{g}_x(\lambda_\ell(W))$. Since $\widetilde{g}_x$ has degree at most b, it suffices for the prover to specify the evaluations of $\widetilde{g}_x(\lambda_\ell(W))$ at $b + 1 = O(\log n)$ points. A direct application of Qqs. (1) and (2) shows that each evaluation can be done in $O(n \log n)$ time, resulting in a total runtime of $O(n \log^2 n)$. However, using now-standard memoization techniques (see e.g. [36, Section 5.1]), it is possible for the prover to in fact perform each of these evaluations in just $O(n)$ time, resulting in a total runtime of $O(n \log n)$. The verifier can run in $O(b) = O(\log n)$ time per stream update, as each stream update x_z only requires the verifier to compute $\chi_{\mathbf{z}}(\mathbf{r})$, and it follows from Eq. (2) that this can be done with $O(b)$ field operations. When interacting with the prover, the verifier first needs to determine the line ℓ through $\mathbf{j}$ and $\mathbf{r}$, which he can do in $O(b) = O(\log n)$ time. To process the prover's reply, he must evaluate the polynomial h at the points t and w; these evaluations can be done in polylog n time. ◀

The above SIP protocol uses very little of the special structure of the INDEX problem. Let us abstract out its salient features, so as to handle the general problem described at the start of this section. First, note the protocol treats the data set given by $(x_1, \ldots, x_n)$ as an implicit description of the polynomial $\widetilde{g}_x$. Second, note that our soundness analysis did not require multilinearity per se, only an upper bound on the total degree of $\widetilde{g}_x$. Finally, note that the specific form of Eqs. (1) and (2) is not crucial either; all we used was that it allows the verifier an easy streaming computation. Thus, we obtain the following generic result.

▶ **Theorem 2.2** (Polynomial Evaluation Protocol). *Suppose an input data stream implicitly describes a v-variate polynomial g of total degree d over a field $\mathbb{F}$, followed by a point $\mathbf{j} \in \mathbb{F}^v$. Suppose this implicit description allows a streaming verifier to evaluate g at a random point $\mathbf{r} \in_R \mathbb{F}^v$ using space S. Then the technique of the protocol in Figure 2 gives a two-round SIP for computing $g(\mathbf{j})$, with the following properties: (1) perfect completeness; (2) soundness error bounded by $d/(|\mathbb{F}| - 1)$; (3) space usage in $O(v \log |\mathbb{F}| + S)$; (4) help cost in $O((d + v) \log |\mathbb{F}|)$.* ◀

We shall refer to the abstract protocol given by Theorem 2.2 as the polynomial evaluation protocol.

3   Constant-Round SIPs for Query Problems

We shall now apply the polynomial evaluation protocol to design SIPs proving the various upper bounds outlined in Section 1.1. The first application is immediate; later applications bring in additional ideas.

3.1   Point Queries.

In the POINTQUERY problem, the input is a stream in the turnstile model, updating an initially-zero vector $\mathbf{x} \in \mathbb{Z}^n$, followed by a query $j \in [n]$. The goal is to output x_j.

▶ **Theorem 3.1.** *Suppose the input to* POINTQUERY *is guaranteed to satisfy* $|x_i| \leq q$ *at end of the data stream, for all entries of* $\mathbf{x}$*, where the bound q is known a priori. Then there is a two-round SIP for* POINTQUERY *with space and help costs in* $O(\log n \log(q + \log n))$.

Proof. Assume WLOG that $n = 2^b$ for an integer b, and use a bijection $u \in [n] \longleftrightarrow \mathbf{u} \in \{0,1\}^b$ as in Theorem 2.1. The vector $\mathbf{x}$ resulting from the updates defines a multilinear polynomial $\widetilde{g}_{\mathbf{x}}(Z_1, \ldots, Z_b)$ by Eq. (1), where $g_{\mathbf{x}}(\mathbf{z}) := x_z$. We can treat $\widetilde{g}_{\mathbf{x}}$ as a polynomial over any field we like, but to solve our problem, we need to tell apart the $2q + 1$ possible values taken on by the entries of $\mathbf{x}$ (recall that q is an upper bound on $\|\mathbf{x}\|_\infty$ at the end of the stream). For this it suffices to have $\operatorname{char}(\mathbb{F}) \geq 2q + 1$.

Applying the polynomial evaluation protocol is now straightforward. The verifier starts with $\mathbf{r} \in_R \mathbb{F}^b$ and $Q = 0$. Upon receiving an update indicating "$x_i \leftarrow x_i + c$," he updates $Q \leftarrow Q + c\chi_{\mathbf{i}}(\mathbf{r})$. The other details are as in Figure 2. The space and communication costs are both in $O(b \log |\mathbb{F}|)$ as before.

To ensure a soundness error of at most $1/3$, we let $|\mathbb{F}| > 3b$ as before. This and the earlier condition on $\operatorname{char}(\mathbb{F})$ can both be satisfied by, e.g., taking $\mathbb{F} = \mathbb{F}_p$, for a prime $p > 3b + 2q$. This translates to cost bounds in $O(\log n \log(q + \log n))$, as claimed.   ◀

3.2   Nearest Neighbor Queries

Consider a "premetric" space [2] $(\mathcal{X}, D)$ given by a finite ground set $\mathcal{X}$ and distance function $D : \mathcal{X} \times \mathcal{X} \to \mathbb{R}^+$ satisfying $D(\mathbf{x}, \mathbf{x}) = 0$ for all $\mathbf{x} \in \mathcal{X}$. Let $B_D(\mathbf{z}, r) = \{\mathbf{x} \in \mathcal{X} : D(\mathbf{x}, \mathbf{z}) \leq r\}$ denote the corresponding ball of radius $r \in \mathbb{R}^+$ centered at $\mathbf{z} \in \mathcal{X}$. In the NEARESTNEIGHBOR problem, the input consists of a stream $\langle \mathbf{x}^{(1)}, \ldots, \mathbf{x}^{(m)} \rangle$ of m points from $\mathcal{X}$, constituting the data set, followed by a query point $\mathbf{z} \in \mathcal{X}$. The goal is to output $\mathbf{x}^\star = \arg\min_{\mathbf{x}^{(i)}} D(\mathbf{x}^{(i)}, \mathbf{z})$, the nearest neighbor of $\mathbf{z}$ in the data set. We shall give highly efficient SIPs for this problem that handle rather general distance functions D. To keep our statements of bounds simple, we shall impose the following structure on $(\mathcal{X}, D)$.

- We assume that $\mathcal{X} = [n]^d$. We think of d as the dimensionality of the data, and $[n]^d$ as a very fine "grid" over the ambient space of possible points.
- For all $\mathbf{x}, \mathbf{y} \in [n]^d$, $D(\mathbf{x}, \mathbf{y}) \leq 1$ is an integer multiple of a small parameter $\varepsilon \geq 1/n^d$.

Overall, this amounts to assuming that our data set has polynomial *spread*: the ratio between the maximum and minimum distance. We proceed to give two SIPs for NEARESTNEIGHBOR. Our basic SIP has cost roughly logarithmic in the stream length and the spread (and therefore

---

[2]   This very general setting, which includes metric spaces as special cases, captures several important distance functions such as the Bregman divergences from information theory and machine learning that satisfy neither symmetry nor the triangle inequality.

linear in d but only logarithmic in n). After we present it, we shall critique it and then give a more sophisticated SIP to handle its faults.

▶ **Theorem 3.2.** *Under the above assumptions on the premetric space $(\mathcal{X}, D)$, the* NEAR-ESTNEIGHBOR *problem has a three-round SIP with cost $O(d \log n \log(m + \log(d \log n)))$.*

Proof. Let $\mathcal{B} = \{B_D(\mathbf{x}, j\varepsilon) : \mathbf{x} \in \mathcal{X}, j \in \mathbb{Z}, 0 \le j \le 1/\varepsilon\}$ be the set of all balls of all radii between 0 and 1 (quantized at granularity ε). By our assumptions on the structure of $(\mathcal{X}, D)$, we have $|\mathcal{B}| \le n^d/\varepsilon \le n^{2d}$. The input stream $\langle \mathbf{x}^{(1)}, \ldots, \mathbf{x}^{(m)} \rangle$ defines a *derived stream*, consisting of updates to a vector $\mathbf{v}$ indexed by the elements of $\mathcal{B}$. We shall denote by $v[\boldsymbol{\beta}]$ the entry of $\mathbf{v}$ indexed by $\boldsymbol{\beta} \in \mathcal{B}$. The derived stream is defined as follows: the token $\mathbf{x}^{(i)}$ increments $v[\boldsymbol{\beta}]$ for every ball $\boldsymbol{\beta}$ that contains $\mathbf{x}^{(i)}$. The verifier runs the POINTQUERY protocol of Theorem 3.1 on this derived stream.

The verifier learns the query point $\mathbf{z}$ at the end of the stream. The prover then supplies a point $\mathbf{y}$ claimed to be a valid nearest neighbor (note that there may be more than one valid answer). To check this claim, it is sufficient for the verifier to check two properties: (1) that $\mathbf{y}$ did appear in the stream, and (2) that the stream contained no point closer to $\mathbf{z}$ than $\mathbf{y}$. The first property holds iff $v[B_D(\mathbf{y}, 0)] \ne 0$. The second property holds iff $v[B_D(\mathbf{z}, D(\mathbf{y}, \mathbf{z}) - \varepsilon)] = 0$. Clearly, these two properties can be checked by two point queries over the derived stream.

Following the protocol of Theorem 3.1, the two point queries (executed in parallel) involve two more rounds between the verifier and the prover, for an overall three-round SIP. Since the entries of $\mathbf{v}$ never exceed m, each POINTQUERY protocol requires space and help costs $O(d \log n \log(m + \log(d \log n)))$. ◀

While the protocol of Theorem 3.2 achieves very small space and help costs, the prover's and verifier's runtimes could be as high as $\Omega(n^d)$, because processing a single stream token $\mathbf{x}^{(i)}$ may require both parties to enumerate all balls containing $\mathbf{x}^{(i)}$. Ultimately, this inefficiency is because the protocol assumes hardly anything about the nature of the distance function D and, as a result, does not get to exploit any structural information about the balls in $\mathcal{B}$.

To rectify this, we shall make the entirely reasonable assumption that the distance function D is "efficiently computable" in the rather mild sense that membership in a ball generated by D can be decided by a short (say, polynomial-length) formula. Accordingly, we shall express our bounds in terms of a parameter that captures this notion of efficient computation.

▶ **Definition 3.3.** Suppose the distance function D on $\mathcal{X}$ satisfies the assumptions for Theorem 3.2. Let $\Phi_D : \mathcal{B} \times \mathcal{X} \to \{0, 1\}$ be the ball membership function for D, i.e., $\Phi_D(B_D(\mathbf{z}, r), \mathbf{x}) = 1 \iff \mathbf{x} \in B_D(\mathbf{z}, r)$. Think of Φ_D as a Boolean function of $(3d \log n)$-bit inputs. We define the *formula size complexity* of D, denoted fsize(D), to be the length of the shortest de Morgan formula for Φ_D.

Since addition and multiplication of b-bit integers can both be computed by Boolean circuits in depth $\log b$ (see, e.g., [31, 37]), they can be computed by Boolean formulae of size poly(b). It follows that for many natural distance functions D, including the Euclidean, Hamming, ℓ_1, and ℓ_∞ metrics (and in fact ℓ_p for all suitably "small" positive p), we have fsize(D) = poly($d, \log n$).

▶ **Theorem 3.4.** *Suppose the premetric space $(\mathcal{X}, D)$ satisfies the assumptions made for Theorem 3.2. Then* NEARESTNEIGHBOR *on $(\mathcal{X}, D)$ has a three-round SIP, whose space and help costs are both at most $O(\text{fsize}(D) \log(m + \text{fsize}(D)))$, in which the verifier processes each*

stream update in time $O(\text{fsize}(D))$, and the prover runs in total time $m \cdot \text{poly}(\text{fsize}(D))$. In particular, if $\text{fsize}(D) = \text{poly}(d, \log n)$, as is the case for many natural distance functions D, then the space and help costs are both $\text{poly}(d, \log m, \log n)$, the verifier runs in time $\text{poly}(d, \log n)$ per stream update, and the prover runs in total time $m \cdot \text{poly}(d, \log n)$.

We defer a proof of Theorem 3.4 to the full version of the paper, but the high level idea that allows us to avoid the high runtimes of the previous protocol is as follows. Essentially, the SIP of Theorem 3.2 ran our polynomial evaluation protocol on a *multilinear* extension of the vector $\mathbf{v}$ defined by the derived stream. That SIP took $\mathbf{v}$ to be a completely arbitrary table of values. As a result, the verifier's computation – evaluating the multilinear extension at a random point – became costly. The honest prover incurred similar costs. A closer examination of the nature of $\mathbf{v}$ reveals that if D is a "reasonable" distance function, then $\mathbf{v}$ itself has plenty of structure. In particular, an appropriate *higher degree* extension of $\mathbf{v}$ can in fact be evaluated much more efficiently (by both the verifier and the prover) than the above multilinear extension.

3.3   Range Counting Queries

Let $\mathcal{U}$ be any data universe and $\mathcal{R} \subseteq 2^{\mathcal{U}}$ a set of *ranges*. In the RANGECOUNT problem, the data stream $\sigma = \langle \mathbf{x}^{(1)}, \ldots, \mathbf{x}^{(m)}, R^* \rangle$ specifies a sequence of universe elements $\mathbf{x}^{(i)} \in \mathcal{U}$, followed by a *query* or *target* range $R^* \in \mathcal{R}$. The goal is to output $|\{i : \mathbf{x}^{(i)} \in R^*\}|$, i.e., the number of elements in the target range that appeared in the stream.

We easily obtain a *two-round* streaming interactive proof for the RANGECOUNT problem with cost bounded by $O\left(\log |\mathcal{R}| \log (|\mathcal{R}|m)\right)$. The verifier simply runs a POINTQUERY on the derived stream σ' defined to have data universe $\mathcal{R}$. σ' is obtained from σ as follows: on each stream update $\mathbf{x}^{(i)} \in \mathcal{U}$, the verifier inserts into σ' one copy of each range $R \in \mathcal{R}$ such that $\mathbf{x}^{(i)} \in R$. The range count problem is equivalent to a POINTQUERY on σ', with the target item being R^*, and we obtain the following theorem.

▶ **Theorem 3.5.** *There is a two-round SIP with $O(\log |\mathcal{R}| \log(|\mathcal{R}|m))$ cost for RANGECOUNT.*

In particular, for spaces of bounded *shatter dimension* ρ, $\log |\mathcal{R}| = \rho \log m = O(\log m)$. The above protocol also implies a *three-round* SIP for the problem of *linear classification*, a core problem in machine learning. Just like the protocol for NEARESTNEIGHBOR invokes a two-round protocol for INDEX, an SIP for linear classification (find a hyperplane that separates red and blue points) verifies that the proposed hyperplane is empty of red points on one side and blue points on the other using the above two-round RANGECOUNT protocol.

The prover and verifier in the protocol of Theorem 3.5 may require time $\Omega(|\mathcal{R}|)$ per stream update. This could be prohibitively large. However, we can obtain savings analogous to Theorem 3.4 if we make a mild "efficient computability" assumption on our ranges. Specifically, suppose there exists a ($\text{poly}(S)$-time uniform) de Morgan formula Φ of length S that takes as input a binary string representing a point $\mathbf{x}^{(i)} \in \mathcal{U}$, as well as the label of a range $R \in \mathcal{R}$ and outputs a bit that is 1 if and only if $\mathbf{x}^{(i)} \in R$. We then obtain the following more practical SIP.

▶ **Theorem 3.6.** *Suppose membership in ranges from $\mathcal{R}$ can be decided by de Morgan formulas of length S as above. Then there is a two-round SIP for RANGECOUNT on $\mathcal{R}$, with costs at most $O(S \log(m + S))$, in which the verifier runs in time $O(S)$ per stream update, and the prover runs in total time $m \cdot \text{poly}(S)$.*

3.4 Median and Selection Queries

We give a three-round SIP for SELECTION, of which MEDIAN is a special case. In the SELECTION problem, defined over data universe $\mathcal{U} = [n]$, the data stream $\sigma = \langle \mathbf{x}^{(1)}, \ldots, \mathbf{x}^{(m)}, \rho \rangle$ is a sequence of elements from $[n]$, followed by a desired rank $\rho \in [m]$. For $i \in [n]$, let $f_i := \{j : \mathbf{x}^{(j)} = i\}$ denote the number of times element i appears in the stream. Given a desired rank $\rho \in [m]$, the goal is to output an element $j \in [n]$ such that

$$\sum_{k < j} f_k < \rho \quad \text{and} \quad \sum_{k > j} f_k \leq m - \rho. \tag{3}$$

MEDIAN is the special case of SELECTION when $\rho = \lfloor m/2 \rfloor$.

Our three-round SIPs for SELECTION essentially work by reducing to the RANGECOUNT problem, but an extra round is required for the prover to send the desired element j to the verifier.

▶ **Theorem 3.7.** *There is a three-round SIP for* SELECTION *with cost at most* $O(\log n \log(m + \log n))$ *in which the verifier runs in time* $\mathrm{poly}(\log n, \log m)$ *per update, and the prover runs in total time* $m \cdot \mathrm{poly}(\log n, \log m)$.

The proof of Theorem 3.7 is deferred to the full version of the paper.

3.5 Pattern Matching Queries

In the pattern matching with wildcards problem, denoted PMW, we are given a stream σ representing text $T = (t_1, \ldots, t_m) \in \{0, 1, *\}^m$ *followed* by a pattern $P = (p_1, \ldots, p_q) \in \{0, 1, *\}^q$. The wildcard symbol $*$ is interpreted as "don't care", and the pattern P is said to occur at location i in t if, for every position j in P, either $p_j = t_{i+j}$ or at least one of p_j and t_{i+j} is the wildcard symbol. The PMW problem is to determine the number of locations at which P occured in T. PATTERNMATCHING refers to the special case where "don't care" symbols are not permitted. We focus on a binary alphabet; a larger alphabet $\mathcal{U}$ can be handled by replacing each character in $\mathcal{U}$ with its binary representation, growing the parameter q by a factor of $\log |\mathcal{U}|$.

Pattern matching, both with and without wildcards, has been extensively studied within the algorithmic literature, with applications ranging from internet search to computational genetics (see e.g. [11, 20] and the references therein). Verifiable protocols for pattern matching enable searching in the cloud, and complements work on searching in encrypted data within the cloud (e.g. [7]). Cormode et al. [13] described and implemented an SIP for PMW that required roughly $\Theta(\log^2 m)$ rounds and had space help costs bounded by $\tilde{\Theta}(\log^2 m)$; Concretely, their implementation required well over 1,000 rounds, even for quite small streams (of length 2^{17}). In stark contrast, our new protocol requires the optimal number of rounds: two.

▶ **Theorem 3.8.** *There is a 2-round SIP for* PMW *with space and help costs at most* $O(q \log(q + m))$, *in which the verifier runs in time* $O(q)$ *per stream update, and the prover runs in total time* $m \cdot \mathrm{poly}(q)$.

The proof of Theorem 3.8 is deferred to the full version of the paper. We remark that the PMW protocol of Theorem 3.8 can be run even if the verifier only knows an upper bound on the length q of the pattern. This is because, for any $q' \leq q$, a pattern $P' \in \{0, 1, *\}^{q'}$ is equivalent to the pattern $P \in \{0, 1, *\}^q$ obtained from P' by concatenating $q - q'$ wildcard symbols to P'.

4   Communication Protocols and Complexity Classes

We now turn to the study of communication complexity classes motivated by a desire to understand streaming interactive proofs (SIPs) from a complexity-theoretic viewpoint. In this section, we lay out the necessary definitions and terminology to rigorously discuss the notions outlined in Section 1.3. In the next section we prove the many parts of Result 1.5.

4.1   Definitions

Communication problems arise naturally out of data stream problems if we suppose Alice holds a prefix of the input stream, and Bob the remaining suffix. The primary goal of such reductions is to obtain space lower bounds on data stream algorithms, so we are free to split the stream at any place we like. For example, most data stream problems in Section 3 are *query problems*, where the input consists of a streamed data set, S, followed by a query, q, to apply to S. In this case, it would be natural to split the input by giving S to Alice and q to Bob. Communication problems that will play an important role in this paper include the index problem INDEX : $\{0,1\}^n \times [n] \to \{0,1\}$ where $[n] := \{1, \dots, n\}$ and INDEX$(x, j) = x_j$, the set-intersection and set-disjointness problems INTER, DISJ : $\{0,1\}^n \times \{0,1\}^n \to \{0,1\}$ where INTER$(x, y) = \neg$DISJ$(x, y) = \bigvee_{i=1}^n (x_i \wedge y_i)$, and the median relation MED : $[n]^m \times [n]^m \to [n]$, where inputs $x, y \in [n]^m \times [n]^m$ are interpreted as two halves of a list of numbers, and the valid output(s) corresponds to the median(s) of the combined list.

4.1.1   Communication Complexity Classes

All our communication models provide random coins and allow two-sided error probability up to a constant; when unspecified, this constant defaults to $1/3$. Given a communication model **C**, we denote the corresponding complexity measure of a problem f by C(f). Following Babai et al. [5], we also denote by **C** the corresponding complexity class, defined as the set of all functions $f : \{0,1\}^n \times \{0,1\}^n \to \{0,1\}$ such that C$(f) = (\log n)^{O(1)}$, i.e., functions that are "easy" in the model **C**.

We let $\mathbf{R}^{[\mathbf{k}, \mathbf{A}]}$ denote the model of randomized communication complexity where Alice and Bob exchange $k \geq 1$ messages in total with Alice sending the first; $\mathbf{R}^{[\mathbf{k}, \mathbf{B}]}$ is similar, except that Bob starts. In the **MA** model, the super-player Merlin, who sees all of the input, broadcasts a message at the start, following which Alice and Bob run a (two-way, arbitrary-round) randomized "verification" protocol. The $\mathbf{MA}^{[\mathbf{k}, \mathbf{A}]}$ and $\mathbf{MA}^{[\mathbf{k}, \mathbf{B}]}$ models are restrictions of **MA** where Merlin speaks only to Bob [3] and the verification protocol following Merlin's single message is restricted to lie in $\mathbf{R}^{[\mathbf{k}, \mathbf{A}]}$ and $\mathbf{R}^{[\mathbf{k}, \mathbf{B}]}$ respectively.

The **MA** model (indeed, its restriction $\mathbf{MA}^{[\mathbf{1}, \mathbf{A}]}$) allows us to simulate 1-round SIPs in an obvious way: Merlin sends Bob the prover's message, and Alice sends Bob the verifier's memory contents after it has processed her prefix of the stream. Notice that the order of the two messages is not important, modulo one crucial consideration: Alice must have a private channel to Bob and the random coins used to generate the message from Alice to Bob must be *hidden coins*, invisible to Merlin but shared between Alice and Bob (which is why we called them "hidden coins" rather than "private coins").

The models $\mathbf{OMA}^{[\mathbf{k}]}$, $\mathbf{OIP}^{[\mathbf{k}]}$, and $\mathbf{OIP}_+^{[\mathbf{k}]}$, for $k \geq 1$, are obtained by extending $\mathbf{MA}^{[\mathbf{1}, \mathbf{A}]}$ to simulate k-round SIP protocols. These communication models work as follows. In each

---

[3] Our definition breaks symmetry between Alice and Bob because our eventual goal is to study online protocols.

case, Alice and Bob first toss some hidden coins. Then, upon receiving the input, two things happen: (1) Merlin and Bob interact for k rounds, with Merlin sending the last message in the interaction, and (2) Alice sends Bob a message, randomized using the hidden coins. After these actions are completed, Bob produces an output in $\{0, 1\}$. The differences between the three series of models are as follows.

- In $\mathbf{OMA}^{[\mathbf{k}]}$, (1) happens before (2) and Bob must interact with Merlin before looking at his input. This is directly analogous to $\mathbf{AM}_{\mathsf{TM}}$; see the discussion in Section 1.2.
- In $\mathbf{OIP}^{[\mathbf{k}]}$, (1) happens before (2) and Bob may look at his input before talking to Merlin.
- Finally, $\mathbf{OIP}^{[\mathbf{k}]}_+$ is like $\mathbf{OIP}^{[\mathbf{k}]}$ except that (2) happens before (1). Thus, Bob's messages may depend on Alice's actual message to Bob, not just on Bob's input and the hidden coins.

In the $\mathbf{AM}$ model, the parties first choose a public random string, then Merlin broadcasts a message to Alice and Bob, who then run a deterministic communication protocol to arrive at a Boolean output. Since Merlin can in fact predict the exact transcript that Alice and Bob will generate following his message, we can assume without loss of generality that after Merlin's message, Alice and Bob output one bit each indicating whether or not they accept Merlin's prediction.

4.1.2   Cost and Value of Protocols

Let $\mathcal{P}$ be a protocol in a model $\mathbf{C}$ involving Merlin. For each input (x, y), $\mathcal{P}$ defines a game between Merlin and Arthur (recall that Alice and Bob together constitute Arthur), wherein Merlin's goal is to make Arthur output 1. We define the *value* $V^{\mathcal{P}}(x, y)$ to be Merlin's probability of winning this game with optimal play. Given a Boolean function f, we say that $\mathcal{P}$ computes f with *soundness error* ε_s and *completeness error* ε_c if, for all x, y we have

$$f(x, y) = 0 \ \Rightarrow \ V^{\mathcal{P}}(x, y) \leq \varepsilon_s, \quad \text{and} \quad f(x, y) = 1 \ \Rightarrow \ V^{\mathcal{P}}(x, y) \geq 1 - \varepsilon_c. \tag{4}$$

When the above holds with $\varepsilon_c = 0$, we say that $\mathcal{P}$ computes f with perfect completeness.

The *verification cost* of $\mathcal{P}$, denoted $\mathrm{vc}(\mathcal{P})$, is the (worst-case) number of bits sent by Alice plus the number of hidden coin tosses; its *help cost* $\mathrm{hc}(\mathcal{P})$ is the number of bits communicated between Merlin and Bob; its communication cost $\mathrm{cc}(\mathcal{P}) = \mathrm{hc}(\mathcal{P}) + \mathrm{vc}(\mathcal{P})$. For a problem f, we define its complexity $\mathrm{C}(f) = \min\{\mathrm{cc}(\mathcal{Q}) : \mathcal{Q} \text{ is a } \mathbf{C} \text{ protocol that solves } f \text{ with } \max\{\varepsilon_s, \varepsilon_c\} \leq 1/3\}$.

4.2   Relations Among Communication Complexity Classes

We prove a number of inclusion and separation results among our "new" communication complexity classes and relate them to previously studied classes. These are summarized in Figure 1, replicated below.

Our results shed light on the landscape of online communication complexity in general.

The simplest online communication model is $\mathbf{R}^{[\mathbf{1}, \mathbf{A}]}$, a.k.a. one-way randomized communication. The result $\mathbf{OIP}^{[\mathbf{1}]} = \mathbf{OIP}^{[\mathbf{1}]}_+ = \mathbf{R}^{[\mathbf{1}, \mathbf{A}]}$ establishes that in the world of online communication, introducing the omniscient but untrusted Merlin into the model is not enough to obtain super-polynomial efficiency improvements, if *interaction* with Merlin is not permitted. The stronger result that $\mathbf{OMA}^{[\mathbf{k}]} = \mathbf{R}^{[\mathbf{1}, \mathbf{A}]}$ for all constants $k > 0$ (this is the full statement of Theorem 5.20) establishes that in the "public coin" setting, the addition of Merlin is not enough to obtain super-polynomial speedups even if interaction with Merlin *is* permitted.
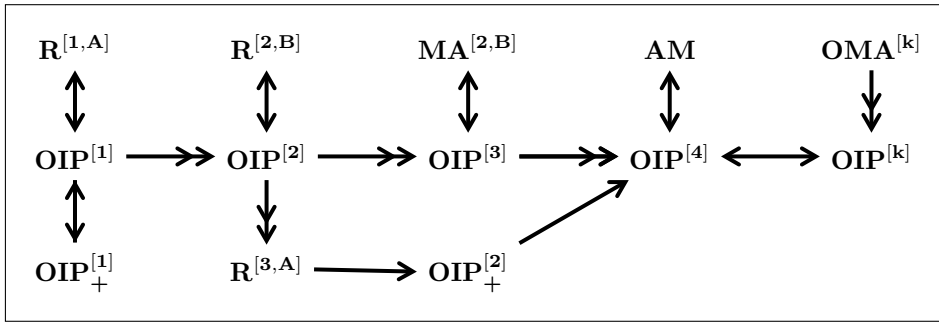
■ **Figure 3** The layout of our communication complexity zoo. An arrow from $\mathbf{C}_1$ to $\mathbf{C}_2$ indicates that $\mathbf{C}_1 \subseteq \mathbf{C}_2$. If the arrow is double-headed, then the inclusion is strict. Within the figure, k is an arbitrary constant larger than 4.

The result that $\mathbf{OIP}^{[2]} = \mathbf{R}^{[2,\mathbf{B}]}$ (see Corollary 5.7) establishes that in the "hidden coin" setting, the addition of Merlin to the communication model *can* yield super-polynomial efficiency improvements, even if only the barest amount of interaction with Merlin is permitted. However, note that $\mathbf{R}^{[2,\mathbf{B}]}$ is the simplest non-online communication model. Thus the combination of hidden coins and a minimal amount of interaction is enough to simulate only the simplest of the non-online communication protocols.

The result that $\mathbf{OIP}^{[4]} = \mathbf{OIP}^{[4]}_+ = \mathbf{AM}$ (see Corollary 5.14) shows that in the "hidden coin" setting, the addition of Merlin to the communication model permits the simulation even of *non-online interactive proofs*, as soon as four rounds of interaction with Merlin are permitted.

This in turn explains the somewhat puzzling result that the $\mathbf{OIP}$ and $\mathbf{OIP}_+$ hierarchies collapse to the fourth level: both Goldwasser–Sipser [16] and Babai–Moran [6] break down in the $\mathbf{OIP}$ and $\mathbf{OIP}_+$ worlds because their transformations *do not preserve online-ness*: they will turn an $\mathbf{OIP}^{[2]}$ protocol into a "public coin" one, but require Merlin to send a message to Alice. However, as soon as four rounds of interaction with Merlin are permitted, even online interactive proofs can simulate non-online ones. At this point, the phenomena of classical interactive proofs kick in, and the hierarchies collapse.

5   A Communication Complexity Zoo

We now study our central communication models $\mathbf{OIP}^{[k]}$ and $\mathbf{OIP}^{[k]}_+$, and prove the web of relationships given in Figure 3. Our results are of two types: (1) establishing separations or collapses between levels of the $\mathbf{OIP}$ and $\mathbf{OIP}_+$ hierarchies, as the case may be, and (2) relating these hierarchies to other previously studied communication complexity classes. We shall first characterize every finite level of the $\mathbf{OIP}$ hierarchy (the vertical bidirectional arrows in Figure 3). Next, in Sections 5.4 and 5.5, we separate the first four levels of the hiearchy (the horizontal double-headed arrows in the figure). Finally, in Section 5.5, we separate the $\mathbf{OIP}$ and $\mathbf{OMA}$ hierarchies.

Throughout Section 5, f will denote an arbitrary communication problem given by a Boolean function $f : \mathcal{X} \times \mathcal{Y} \to \{0,1\}$, and n will parametrize its "instance size" up to a constant factor, i.e., we will have $\log|\mathcal{X}| + \log|\mathcal{Y}| = \Theta(n)$. We shall use big-$O$ and big-Ω notation to hide constants independent of f, $|\mathcal{X}|$ and $|\mathcal{Y}|$. We shall use the term "ordinary protocol" to mean a randomized communication protocol involving Alice and Bob alone (and no Merlin).

The first level of the hierarchy is easy to characterize.

▶ **Theorem 5.1.** *We have* $\mathbf{OMA}^{[1]} = \mathbf{OIP}^{[1]} = \mathbf{OIP}^{[1]}_{+} = \mathbf{MA}^{[1,\mathbf{A}]} = \mathbf{R}^{[1,\mathbf{A}]}$.

Proof. The definitions immediately show that the first four classes are identical (syntactically) and include $\mathbf{R}^{[1,\mathbf{A}]}$, because one can always choose to ignore Merlin. The reverse inclusion $\mathbf{MA}^{[1,\mathbf{A}]} \subseteq \mathbf{R}^{[1,\mathbf{A}]}$ follows from previous work: Chakrabarti et al. [9] show that for all f we have $\mathrm{R}^{[1,A]}(f) = O\big(\mathrm{MA}^{[1,A]}(f)^2\big)$. ◀

5.1 A Characterization of $\mathrm{OIP}^{[2]}$

The main goal of this subsection is to prove that $\mathbf{OIP}^{[2]} = \mathbf{R}^{[2,\mathbf{B}]}$. We start with the following communication analog of Theorem 2.2, which was proven in a streaming setting.

▶ **Lemma 5.2** (Polynomial Evaluation Protocol, Communication Version). *Suppose Alice holds a v-variate polynomial g of total degree d over a field $\mathbb{F}$, and Bob holds a point $\mathbf{j} \in \mathbb{F}^v$. Assume $|\mathbb{F}| > 4d$. Then there is an $\mathbf{OIP}^{[2]}$ protocol with communication cost $O((v + d) \cdot \log |\mathbb{F}|)$ for evaluating $g(\mathbf{j})$. In particular, $\mathrm{OIP}^{[2]}(\textsc{index}) = O(\log n \log \log n)$, so that $\textsc{index} \in \mathbf{OIP}^{[2]}$.*

Proof. Using the notation established in the description of the polynomial agreement protocol of Section 2.1, we simply note that the hidden coins shared between Alice and Bob determine $\mathbf{r}$. Bob can send Merlin (the canonical representation of) the line ℓ that passes through $\mathbf{j}$ and $\mathbf{r}$ without having to hear from Alice, since ℓ is determined entirely by $\mathbf{r}$ and Bob's input $\mathbf{j}$. Merlin can send Bob the polynomial $h(W)$ claimed to equal g restricted to the line ℓ, and Alice can send Bob $g(\mathbf{r})$ within the stated cost bounds. Bob performs the same check as the verifier in Theorem 2.2, and the completeness and soundness analysis is unchanged. ◀

The just-proved fact that $\textsc{index} \in \mathbf{OIP}^{[2]}$ is striking: combined with the well-known lower bound $\mathbf{R}^{[1,\mathbf{A}]}(\textsc{index}) = \Omega(n)$, it shows that introducing Merlin into the picture while keeping the one-way restriction on the Alice/Bob communication lowers cost exponentially. It is now natural to ask whether $\mathbf{OIP}^{[2]}$ allows such exponential savings for harder problems, such as $\textsc{disj}$. Our next result – a lower bound on $\mathbf{OIP}^{[2]}$ complexity – implies that it does not.

▶ **Theorem 5.3.** *Let $\mathcal{P}$ be an $\mathbf{OIP}^{[2]}$ protocol computing f. Then $\mathrm{hc}(\mathcal{P})\,\mathrm{vc}(\mathcal{P}) = \Omega(\mathrm{R}^{[2,B]}(f))$. In particular, $\mathrm{OIP}^{[2]}(f) = \Omega\big(\mathrm{R}^{[2,B]}(f)^{1/2}\big)$, which implies $\mathbf{OIP}^{[2]} \subseteq \mathbf{R}^{[2,\mathbf{B}]}$.*

Proof. After appropriate parallel repetition, we may assume that the soundness and completeness errors of $\mathcal{P}$ at most $1/12$ each. In general, $\mathcal{P}$ takes the following shape: (1) hidden coins are tossed, generating random string r according to distribution $\mathcal{D}$; (2) Bob sends Merlin a message $m_B = m_B(y, r)$; (3) Merlin responds with a message $m_M = m_M(x, y, m_B)$; (4) Alice sends Bob a message $m_A = m_A(x, r)$; (5) Bob outputs a bit given by a function $\mathrm{out}^P(y, m_M, m_A)$.

Let $\mathcal{D}_m$ be $\mathcal{D}$ conditioned on the event $\{m_B(y, r) = m\}$. Note that the distribution $\mathcal{D}_m$ depends on *both* y and m. Since Bob knows y, Bob can determine the distribution $\mathcal{D}_m$ for any value of m (this is not, however, true for Alice, because Alice does not know y).

With this notational setup, we now describe (in Figure 4) a two-message ordinary protocol $\mathcal{Q}$ that we claim computes f.

To analyze this protocol, let us first define the *weight* $W_{x,y}(\overline{m})$ of a Bob-message $\overline{m}$ to be the probability that Merlin, playing optimally after receiving $\overline{m}$, convinces Bob to output 1. That is,

$$W_{x,y}(\overline{m}) = \max_{m_M} \Pr_{r \sim \mathcal{D}_{\overline{m}}} \left[\mathrm{out}^{\mathcal{P}}(y, m_M, m_A(x, r)) = 1 \right]. \tag{5}$$

---

1. Bob samples $\overline{r} \sim \mathcal{D}$, computes $\overline{m} = m_B(y, \overline{r})$, then sends Alice i.i.d. samples $r^{(1)}, \ldots, r^{(h)} \sim \mathcal{D}_{\overline{m}}$, where $h = 36(\mathrm{hc}(\mathcal{P}) + 4)$.
2. Alice sends Bob $m_A(x, r^{(1)}), \ldots, m_A(x, r^{(h)})$.
3. Bob outputs 1 iff $\exists m_M : |\{i \in [h] : \mathrm{out}^{\mathcal{P}}(y, m_M, m_A(x, r^{(i)})) = 1\}| > h/2$.

■ **Figure 4** The $\mathbf{R}^{[2,\mathbf{B}]}$ protocol $\mathcal{Q}$, which simulates the $\mathbf{OIP}^{[2]}$ protocol $\mathcal{P}$.

Then, with $\overline{m} \sim m_B(y, \mathcal{D})$, the expected weight $\mathbb{E}_{\overline{m}}[W_{x,y}(\overline{m})]$ is at least $11/12$ when $f(x,y) = 1$ and at most $1/12$ when $f(x,y) = 0$.

Correctness on 1-inputs: Fix $(x,y) \in f^{-1}(1)$. We shall proceed assuming that the specific Bob-message $\overline{m}$ chosen in Step 1 of $\mathcal{Q}$ satisfies $W_{x,y}(\overline{m}) > 2/3 = 1 - 4(1/12)$; by Markov's inequality, this fails to happen with probability at most $1/4$. Studying Eq. (5) tell us that there exists a specific Merlin-message m_M^* such that $\Pr_r[\mathrm{out}^{\mathcal{P}}(y, m_M^*, m_A(x,r)) = 1] > 2/3$. Therefore, according to the strategy in Steps 2 and 3, the size of the set $\{i \in [h] : \mathrm{out}^{\mathcal{P}}(y, m_M^*, m_A(x, r^{(i)})) = 1\}$ is a sum of h i.i.d. indicators and exceeds $2h/3$ in expectation. By standard Chernoff bounds (e.g., [28, Theorem 4.4]), the probability that Bob outputs 0 is $2^{-\Omega(h)}$. Thus, overall, the probability that $\mathcal{Q}$ outputs 0 on input (x,y) is at most $1/4 + 2^{-\Omega(h)} < 1/3$.

Correctness on 0-inputs: Fix $(x,y) \in f^{-1}(0)$. We shall proceed assuming that the specific Bob-message $\overline{m}$ chosen in Step 1 of $\mathcal{Q}$ satisfies $W_{x,y}(\overline{m}) < 1/3$; by Markov's inequality, this fails to happen with probability at most $1/4$. For each specific Merlin-message m_M, define

$$\mathrm{size}(m_M) = \left|\{i \in [h] : \mathrm{out}^{\mathcal{P}}(y, m_M, m_A(x, r^{(i)})) = 1\}\right|.$$

Then $\mathrm{size}(m_M)$ is a sum of h i.i.d. indicators and has expectation below $h/3$. By standard Chernoff bounds, $\Pr[\mathrm{size}(m_M) > h/2] \leq e^{-h/36}$. By a union bound over all possible Merlin-messages m_M, the probability that Bob outputs 1 is at most $2^{\mathrm{hc}(\mathcal{P})} e^{-h/36} < 2^{-4}$, using our choice of h. Adding in the $1/4$ from our Markov argument earlier, the overall probability that $\mathcal{Q}$ outputs 1 on input (x,y) is at most $1/4 + 2^{-4} < 1/3$.

Communication Cost: By definition of the $\mathbf{OIP}^{[2]}$ model, we have $|r| \leq \mathrm{vc}(\mathcal{P})$ and $|m_A| \leq \mathrm{vc}(\mathcal{P})$. Thus, each of the two messages in $\mathcal{Q}$ costs at most $h \cdot \mathrm{vc}(\mathcal{P}) = O(\mathrm{hc}(\mathcal{P}) \mathrm{vc}(\mathcal{P}))$ bits. ◀

The above proof exploits a key property of $\mathbf{OIP}^{[2]}$ protocols: Bob can sample from the conditional distribution $\mathcal{D}_{\overline{m}}$. This is possible because $m_B = m_B(y, r)$ is independent of Alice's message m_A, a property not satisfied in the stronger $\mathbf{OIP}_+^{[2]}$ model. This explains why Theorem 5.3 does not apply to $\mathbf{OIP}_+^{[2]}$, and indeed we shall later give an exponential separation between $\mathbf{OIP}^{[2]}$ and $\mathbf{OIP}_+^{[2]}$ in Corollary 5.19.

Theorem 5.3 implies a number of lower bounds for specific problems. We begin with DISJ.

▶ **Corollary 5.4.** *We have $\Omega(n^{1/2}) \leq \mathrm{OIP}^{[2]}(\text{DISJ}) \leq O(n^{1/2} \log n)$. In particular, DISJ $\notin$ $\mathbf{OIP}^{[2]}$.*

Proof. For the lower bound, we combine Theorem 5.3 with the fact that $\mathrm{R}^{[2,B]}(\text{DISJ}) \geq \mathrm{R}(\text{DISJ}) = \Omega(n)$, the last step being a celebrated lower bound [21]. The upper bound follows from the Aaronson–Wigderson protocol [2] for DISJ, which is in fact an $\mathbf{MA}^{[1,\mathbf{A}]}$ protocol. ◀

We remark that we may replace DISJ in Corollary 5.4 with IP2, the "inner product mod 2" function. Indeed, the Aaronson–Wigderson protocol also applies to IP2, and $\mathrm{R}(\text{IP2}) = \Omega(n)$.

Recall that MED is a relation on inputs in $[n]^m \times [n]^m$. Corollary 5.5 below establishes a lower bound of $\Omega(m^{1/4})$ on the cost of any $\mathbf{OIP^{[2]}}$ protocol for MED (the proof is deferred to the full version of the paper). This justifies our use of three rounds in the polylogarithmic cost SIP for MEDIAN we gave in Theorem 3.7, as it implies that any 2-round SIP for median based on known techniques must have polynomial cost.

▶ **Corollary 5.5.** *We have* $\Omega(m^{1/4}) \leq \mathrm{OIP}^{[2]}(\text{MED}) \leq O(m^{1/2} \log^{3/2} n)$.

We have now seen that up to polynomial (specifically, quadratic) blowup, $\mathbf{OIP^{[2]}}$ is no more powerful than ordinary $\mathbf{R^{[2,B]}}$. We now show that up to another quadratic blowup this is in fact a characterization.

▶ **Theorem 5.6.** *For all f, we have* $\mathrm{OIP}^{[2]}(f) = O\big(\mathrm{R}^{[2,B]}(f)^2\big)$. *In particular,* $\mathbf{OIP^{[2]}} \supseteq \mathbf{R^{[2,B]}}$.

Proof. Let $\mathcal{Q}$ be an $\mathbf{R^{[2,B]}}$ protocol for f with cost C and error at most $1/6$. Assume WLOG that $C \geq 5$ and that each of the two messages in $\mathcal{Q}$ is a string in $\{0,1\}^C$. We shall treat Alice's messages as elements of the field $\mathbb{F} = \mathbb{F}_{2^C}$ via an agreed-upon bijection.

We design an $\mathbf{OIP^{[2]}}$ protocol $\mathcal{P}$ for f, based on $\mathcal{Q}$. Given an input (x, y), $\mathcal{P}$ begins by choosing a (hidden) random string r shared between Alice and Bob exactly as $\mathcal{Q}$ would have. From now on, think of x, y, r as fixed. This then fixes a message m_B that Bob would have sent Alice in $\mathcal{Q}$, as well as a function $m_A: \{0,1\}^C \to \mathbb{F}$ specifying Alice's response to each Bob-message. Let $\widetilde{m}_A(Z_1, \ldots, Z_C) \in \mathbb{F}[Z_1, \ldots, Z_C]$ be the multilinear extension of this function m_A. In $\mathcal{P}$, Alice needs to send a message to Bob that allows him to determine $m_A(m_B) = \widetilde{m}_A(m_B)$ with Merlin's help. This is an instance of polynomial evaluation, so we solve it by applying the $\mathbf{OIP^{[2]}}$ polynomial evaluation protocol (PEP) from Lemma 5.2.

The polynomial $\widetilde{m}_A$ is C-variate and has total degree C. Therefore, PEP has communication cost $O(C \log |\mathbb{F}|) = O(C^2)$, as does $\mathcal{P}$. Next, PEP has perfect completeness, so an honest Merlin can cause $\mathcal{P}$ to output 1 whenever the choice of r would have caused $\mathcal{Q}$ to output 1. Finally, PEP has soundness error at most $C/(|\mathbb{F}| - 1) = C/(2^C - 1) < 1/6$, so a dishonest Merlin can cause $\mathcal{P}$ to differ in output from $\mathcal{Q}$ with probability at most $1/6$. Using the error bound of $1/6$ on $\mathcal{Q}$, we conclude that $\mathcal{P}$ has completeness error at most $1/6$ and soundness error at most $1/6 + 1/6 = 1/3$. ◀

▶ **Corollary 5.7.** *For all f, we have* $\Omega\big(\mathrm{R}^{[2,B]}(f)^{1/2}\big) \leq \mathrm{OIP}^{[2]}(f) \leq O\big(\mathrm{R}^{[2,B]}(f)^2\big)$. *Thus,* $\mathbf{OIP^{[2]}} = \mathbf{R^{[2,B]}}$.

Proof. Combine Theorems 5.3 and 5.6. ◀

5.2 A Characterization of $\mathrm{OIP}^{[3]}$

The main goal of this subsection is to prove that $\mathbf{OIP^{[3]}} = \mathbf{MA^{[2,B]}}$. Theorem 5.8 below gives a lower bound that builds on the argument in Theorem 5.3 (the proof is deferred to the full version of the paper). Just as before, we can then derive a lower bound for the specific problem DISJ.

▶ **Theorem 5.8.** *Let $\mathcal{P}$ be an $\mathbf{OIP^{[3]}}$ protocol computing f. Then there is an $\mathbf{MA^{[2,B]}}$ protocol $\mathcal{Q}$ computing f with* $\mathrm{hc}(\mathcal{Q}) \leq \mathrm{hc}(\mathcal{P})$ *and* $\mathrm{vc}(\mathcal{Q}) = O(\mathrm{hc}(\mathcal{P}) \mathrm{vc}(\mathcal{P}))$. *In particular,* $\mathrm{OIP}^{[3]}(f) = \Omega\big(\mathrm{MA}^{[2,B]}(f)^{1/2}\big)$, *which implies* $\mathbf{OIP^{[3]}} \subseteq \mathbf{MA^{[2,B]}}$.

▶ **Corollary 5.9.** *We have* $\Omega(n^{1/3}) \leq \text{OIP}^{[3]}(\text{DISJ}) \leq O(n^{1/3} \log n)$. *In particular,* DISJ $\notin$ **OIP**$^{[3]}$.

Proof. Klauck [22] proved that $\text{MA}(\text{DISJ}) = \Omega(n^{1/2})$. Applying Theorem 5.8 to this result gives the non-tight bound $\text{OIP}^{[3]}(\text{DISJ}) = \Omega(n^{1/4})$. But we observe that Klauck's proof shows something stronger: namely, if an **MA** protocol $\mathcal{Q}$ computes DISJ, then $\text{hc}(\mathcal{Q})\,\text{vc}(\mathcal{Q}) = \Omega(n)$. Combining Theorem 5.8 with *this* result, we conclude that if an **OIP**$^{[3]}$ protocol $\mathcal{P}$ computes DISJ, then $\text{hc}(\mathcal{P})^2\,\text{vc}(\mathcal{P}) = \Omega(n)$, and therefore $\text{hc}(\mathcal{P}) + \text{vc}(\mathcal{P}) = \Omega(n^{1/3})$.

For the upper bound, we note that Aaronson and Wigderson [2] also gave an online **MAMA** protocol for DISJ of cost $O(n^{1/3} \log n)$. Every online **MAMA** protocol admits a simulation in **OIP**$^{[3]}$. ◀

As with Corollary 5.4, we may replace DISJ in the above result with IP2. Indeed, Klauck's result [22] implies that $\text{MA}(\text{IP2}) = \Omega(n^{1/2})$, and Aaronson and Wigderson's **MAMA** protocol also applies to IP2.

As we did for the second level in the **OIP** hierarchy, we give an upper bound that applies to the third level and gives a characterization that is tight up to a quadratic blowup.

▶ **Theorem 5.10.** *For all f, we have* $\text{OIP}^{[3]}(f) = O\big(\text{MA}^{[2,B]}(f)^2\big)$. *In particular,* **OIP**$^{[3]} \supseteq$ **MA**$^{[2,B]}$.

Proof sketch. We build on the argument in Theorem 5.6 exactly as the proof of Theorem 5.8 builds on Theorem 5.3. Given an **MA**$^{[2,B]}$ protocol $\mathcal{Q}$ of cost C, the verification strategy used by Alice and Bob in $\mathcal{Q}$ is an **R**$^{[2,B]}$ protocol of cost C, which we can replace with an **OIP**$^{[2]}$ protocol of cost $O(C^2)$, by Theorem 5.6. After this replacement we have an **OIP**$^{[3]}$ protocol. The remaining analysis is routine. ◀

▶ **Corollary 5.11.** *For all f,* $\Omega\big(\text{MA}^{[2,B]}(f)^{1/2}\big) \leq \text{OIP}^{[3]}(f) \leq O\big(\text{MA}^{[2,B]}(f)^2\big)$. *Thus,* **OIP**$^{[3]} = $ **MA**$^{[2,B]}$.

Proof. Combine Theorems 5.8 and 5.10. ◀

5.3 A Characterization of $\text{OIP}^{[4]}$ and Beyond

The fourth level of the **OIP** hierarchy turns out to have surprising power. It can capture all of **AM**, a model that lies at the frontier of our current understanding of communication complexity classes in the sense that we do not know any nontrivial **AM** lower bounds. Thanks to this surprising power, we can show that all constant-height levels of the **OIP** hierarchy collapse to the fourth level.

▶ **Theorem 5.12.** *For all f, we have* $\text{OIP}^{[4]}(f) = O(\text{AM}(f) \log \text{AM}(f))$. *In particular,* **OIP**$^{[4]} \supseteq$ **AM**.

Proof. Suppose $\text{AM}(f) = C$. WLOG, there is a protocol $\mathcal{Q}$ for f with the following shape: Bob tosses coins to generate a random string r and sends it to Merlin, who responds with a message m, where $|r| + |m| \leq C$. Bob then sends (r, m) to Alice, who responds with a single bit, after which Bob announces the output.

The interaction between Bob and Alice is an **R**$^{[2,B]}$ protocol (in fact, it is deterministic) of cost C. Theorem 5.6 shows that it can be replaced with an **OIP**$^{[2]}$ protocol of cost $O(C^2)$. Performing this replacement gives us an **OIP**$^{[4]}$ protocol for f. The cost bound can be improved to $O(C \log C)$ by revisiting the analysis of the polynomial evaluation protocol used to prove Theorem 5.6 and using the fact that Alice's message in $\mathcal{Q}$ is just a single bit. ◀

▶ **Theorem 5.13.** *For each $k > 0$, there exists a constant $c_k > 0$ such that for all f, $\mathrm{OIP}_+^{[k]}(f) \geq \Omega\big(\mathrm{AM}(f)^{c_k}\big)$. In particular, for every constant k, we have $\mathbf{OIP}_+^{[\mathbf{k}]} \subseteq \mathbf{AM}$.*

Proof. Let $C = \mathrm{OIP}_+^{[k]}(f)$ and let $\mathcal{P}$ be an $\mathbf{OIP}_+^{[\mathbf{k}]}$ protocol with cost C that computes f. By definition, $\mathcal{P}$ uses a hidden random string and Merlin learns about this string only indirectly, from Bob's computed messages. We apply the Goldwasser–Sipser set lower bound technique [16] to convert $\mathcal{P}$ into a protocol where all random coins are directly revealed. Specifically, we can convert $\mathcal{P}$ into an $\mathbf{AMAM} \cdots \mathbf{AM}$ protocol $\mathcal{Q}'$, where $k + 3$ messages are sent in total: Merlin's messages are broadcast and after his final message Alice sends a message to Bob, who announces the output. We have $\mathrm{cc}(\mathcal{Q}') = O(C^{a_k})$ for some constant $a_k \geq 1$.

We apply Babai and Moran's round elimination techniques [6] to turn $\mathcal{Q}'$ into a standard $\mathbf{AM}$ protocol $\mathcal{Q}$ of cost at most $O(\mathrm{cc}(\mathcal{Q}')^{b_k})$ for some constant $b_k \geq 1$. The result follows by taking $c_k = 1/(a_k b_k)$. ◀

▶ **Corollary 5.14.** *For all f, $\Omega\big(\mathrm{AM}(f)^{c_4}\big) \leq \mathrm{OIP}^{[4]}(f) \leq O(\mathrm{AM}(f) \log \mathrm{AM}(f))$, where c_4 is the constant from Theorem 5.13. In particular, $\mathbf{OIP}^{[\mathbf{4}]} = \mathbf{AM}$, and in fact $\mathbf{OIP}^{[\mathbf{k}]} = \mathbf{AM}$ for every constant $k \geq 4$.*

Proof. Combine Theorems 5.13 and 5.12, noting that $\mathbf{OIP}^{[\mathbf{k}]} \subseteq \mathbf{OIP}_+^{[\mathbf{k}]}$ for every $k \geq 4$. ◀

Here is an interesting point worth contemplating. On the one hand, our transformations in the proof of Theorem 5.13 perform round reduction at the expense of destroying online-ness: the final protocol $\mathcal{Q}$ is no longer online, i.e., we cannot require communications to go to Bob alone. On the other hand, the transformation in the proof of Theorem 5.12 "restores" onlineness at only a "slight" expense of requiring four rounds, whereas $\mathbf{AM}$ uses only two. Overall, we have a collapse of the $\mathbf{OIP}$ hierarchy to its fourth level.

We have also noted earlier that we (regretfully) do not yet know how to place a concrete problem outside $\mathbf{OIP}_+^{[\mathbf{2}]}$. Nevertheless, Theorems 5.12 and 5.13 together establish a weakness of $\mathbf{OIP}_+^{[\mathbf{2}]}$: up to polynomial factors this model is no more powerful than $\mathbf{OIP}^{[\mathbf{4}]}$.

5.4   Exponential Separations in Our Complexity Zoo

Among the first four levels of the $\mathbf{OIP}$ hierarchy, we can now show that every pair of adjacent levels is exponentially separated. The next three results make this precise. Recall that INTER $= \neg$DISJ is the set intersection problem.

▶ **Theorem 5.15.** *We have $\mathrm{OIP}^{[1]}(\text{INDEX}) = \Omega(n^{1/2})$ whereas $\mathrm{OIP}^{[2]}(\text{INDEX}) = O(\log n \log \log n)$.*

Proof. Combine Theorems 5.1 and Lemma 5.2, and then the known results that $\mathrm{MA}^{[1,A]}(f) = \Omega\big(\mathrm{R}^{[1,A]}(f)^{1/2}\big)$ for all f [9] (see also Theorem 5.20 in Section 5.5), and that $\mathrm{R}^{[1,A]}(\text{INDEX}) = \Omega(n)$ [3]. ◀

▶ **Theorem 5.16.** *We have $\mathrm{OIP}^{[2]}(\text{INTER}) = \Omega(n^{1/2})$ whereas $\mathrm{OIP}^{[3]}(\text{INTER}) = O(\log^2 n)$.*

Proof. For the lower bound, use $\mathrm{R}^{[2,B]}(\text{INTER}) \geq \mathrm{R}(\text{INTER}) = \mathrm{R}(\text{DISJ}) = \Omega(n)$ and then apply Theorem 5.3.

For the upper bound, note that INTER has a nondeterministic protocol with cost $O(\log n)$, wherein Alice and Bob guess an element in the intersection of their respective sets and they verify membership. In particular this gives $\mathrm{MA}^{[2,B]}(\text{INTER}) = O(\log n)$; in fact, Bob need not send anything to Alice in the $\mathbf{MA}^{[\mathbf{2},\mathbf{B}]}$ protocol. Now apply Theorem 5.10. ◀

While we do not know of a *total* Boolean function that separates $\mathbf{OIP^{[3]}}$ from $\mathbf{OIP^{[4]}}$, we do know of a *partial* Boolean function whose $\mathbf{OIP^{[3]}}$ communication complexity is exponentially larger than its $\mathbf{OIP^{[4]}}$ communication complexity. Specifically, Klauck [23, Corollary 3] gives a promise problem he calls PAppMP which has *Quantum* Merlin-Arthur (**QMA**) communication complexity $\Omega(n^{1/6})$ and **AM** communication complexity $O(\log n)$. Since Theorem 5.8 shows that any $\mathbf{OIP^{[3]}}$ protocol can be transformed into an equivalent $\mathbf{MA^{[2,B]}}$ protocol with a quadratic blowup in cost, and $\mathbf{MA^{[2,B]}}$ protocols are simply restricted versions of QMA protocols, Klauck's lower bound on the QMA cost of PAppMP implies that $\mathrm{OIP}^{[3]}(\mathrm{PAppMP}) = \Omega(n^{1/12})$.

Meanwhile, Theorem 5.12 shows that any **AM** communication protocol can be transformed into an equivalent $\mathbf{OIP^{[4]}}$ protocol with a logarithmic blowup in costs. Thus, Klauck's upper bound on the **AM** communication complexity of PAppMP implies that $\mathrm{OIP}^{[4]}(\mathrm{PAppMP}) = O(\log n \log \log n)$.

▶ **Theorem 5.17.** *We have* $\mathrm{OIP}^{[3]}(\mathrm{PAppMP}) = \Omega(n^{1/12})$ *whereas* $\mathrm{OIP}^{[4]}(\mathrm{PAppMP}) = O(\log n \log \log n)$.

Next, we show that, up to polynomial factors, $\mathbf{OIP_+^{[2]}}$ is at least as powerful as $\mathbf{R^{[3,A]}}$, the class of *three-message* randomized communication protocols in which Alice speaks first. This will enable us to exhibit an explicit function f on domain $\{-1, 1\}^n \times \{-1, 1\}^n$ such that $\mathrm{OIP}^{[2]}(f) = \Omega(\sqrt{n/\log n})$, while $\mathrm{OIP}_+^{[2]}(f) = O(\log^2 n)$.

▶ **Theorem 5.18.** *For all f, we have* $\mathrm{OIP}_+^{[2]}(f) = O\big(\mathrm{R}^{[3,A]}(f)^2\big)$.

Proof. Let $\mathcal{Q}$ be any three-message randomized communication protocol of cost C, with Alice speaking first. We show how to convert $\mathcal{Q}$ into an $\mathbf{OIP_+^{[2]}}$ protocol $\mathcal{P}$ of cost $O(C^2)$.

We think of $\mathcal{Q}$ as consisting of one message $m_A^{(1)}$ from Alice to Bob, followed by a *two-message* communication protocol $\mathcal{Q}'$ in which Bob speaks first. Theorem 5.6 shows how to transform $\mathcal{Q}'$ into an equivalent $\mathbf{OIP^{[2]}}$ protocol $\mathcal{P}'$ of cost $O(C^2)$ (note this $\mathbf{OIP^{[2]}}$ protocol *depends* on $m_A^{(1)}$).

Thus, we obtain an $\mathbf{OIP_+^{[2]}}$ protocol $\mathcal{P}$ as follows. Alice's message to Bob in $\mathcal{P}$ consists of two parts. The first specifies $m_A^{(1)}$, and the second is the message she would have sent to Bob in $\mathcal{P}'$. Bob, who learns $m_A^{(1)}$ from the first part of Alice's message, now knows what $\mathbf{OIP^{[2]}}$ protocol $\mathcal{P}'$ to execute, and simply behaves the same as he would in $\mathcal{P}'$. ◀

Exponential separations between $\mathbf{R^{[3,A]}}$ and $\mathbf{R^{[2,B]}}$ are known. In particular, consider the k-step (bipartite) pointer jumping function PJ_k, which interprets each of Alice and Bob's inputs as a list of $N = \Theta(n/\log n)$ *pointers*, a pointer being a $(\log N)$-bit integer. Each pointer in a player's list is interpreted as pointing to (i.e., indexing) a pointer in the other player's list. The goal is to follow these pointers, starting at the first pointer in Alice's list, and output the kth pointer encountered. For example, if Alice's input is $x = (00, 01, 10, 00)$ and Bob's input is $y = (01, 10, 11, 00)$, then $\mathrm{PJ}_1(x, y) = 01$, $\mathrm{PJ}_2(x, y) = 01$, $\mathrm{PJ}_3(x, y) = 10$, and so on. To turn PJ_k into a Boolean function BPJ_k, we take the parity of the $(\log N)$-bit output of PJ_k.

▶ **Corollary 5.19.** *We have* $\mathrm{OIP}^{[2]}(\mathrm{BPJ}_2) = \Omega(\sqrt{n/\log n})$, *while* $\mathrm{OIP}_+^{[2]}(\mathrm{BPJ}_2) = O(\log^2 n)$.

Proof. Nisan and Wigderson [30] showed that $\mathrm{R}^{[k,B]}(\mathrm{BPJ}_k) = \Omega(N/k^2 - k \log N)$. In particular, any two-message randomized communication protocol in which Bob speaks first has cost $\Omega(N)$. Hence, Theorem 5.3 implies that $\mathrm{OIP}^{[2]}(\mathrm{BPJ}_2) = \Omega(\sqrt{n/\log n})$.

To prove the upper bound on $\mathrm{OIP}_+^{[2]}(\text{BPJ}_2)$, note that there is a trivial three-message protocol for PJ_2 (and hence for BPJ_2) of cost $O(\log n)$ in which Alice speaks first. Now apply Theorem 5.18. ◀

5.5 An Exponential Separation Between $\mathrm{OIP}^{[2]}$ and $\mathrm{OMA}^{[k]}$

Theorem 5.20 establishes that for any function f, $\mathrm{OMA}^{[2k]}(f) = \Omega\big(\mathrm{R}^{[1,A]}(f)^{1/(k+1)}\big)$. An essentially identical lower bound was proven by Klauck and Prakash for a closely related (though not identical) communication model; the full version of the paper provides a detailed proof for completeness, and in the process identifies the crucial details of the communication model that enable the lower bound to hold.

▶ **Theorem 5.20.** *For any function f and constant k, $\mathrm{OMA}^{[2k]}(f) = \Omega\big(\mathrm{R}^{[1,A]}(f)^{1/(k+1)}\big)$.*

The main property of the **$\mathrm{OMA}^{[k]}$** communication model exploited in our proof of Theorem 5.20 is the following: in any **$\mathrm{OMA}^{[k]}$** protocol $\mathcal{P}$, for all $i \leq k$, Alice can determine Bob's ith message to Merlin in $\mathcal{P}$ on her own. In particular, the same lower bound would apply to any variant of online Arthur-Merlin communication models in which Bob's messages to Merlin must be independent of his input y. This is the intuitive reason why the **$\mathrm{OIP}^{[2]}$** model is exponentially more powerful than the **$\mathrm{OMA}^{[k]}$** model for any constant k: in the **$\mathrm{OIP}^{[2]}$** model, Bob's message to Merlin may depend on his input y, while this is not allowed in the **$\mathrm{OMA}^{[k]}$** model.

Combining Theorem 5.20 with Lemma 5.2, which says that $\mathrm{OIP}^{[2]}(\text{INDEX}) = O(\log n \log \log n)$, we obtain an exponential separation between **$\mathrm{OIP}^{[2]}$** and **$\mathrm{OMA}^{[k]}$** for any constant $k > 0$.

▶ **Corollary 5.21.** *For every constant $k > 0$, we have $\mathbf{OIP}^{[2]} \not\subseteq \mathbf{OMA}^{[k]}$.* ◀

6   Conclusion

Our primary objects of study in this paper were constant-round interactive protocols for verifying outsourced streaming computations. Our main algorithmic contributions were to give constant-round streaming interactive proofs for a large class of "query" problems. Our protocols are exponentially more efficient than what was believed possible based on prior work, and demonstrate that in the streaming setting, "hidden" coins are exponentially more powerful than public coins.

We also introduced new "online" communication hierarchies, **OIP_+** and **OIP**, which can be seen as restricted variants of the standard Arthur-Merlin communication model. The flow of information in the **OIP_+** and **OIP** models is severely restricted (neither Bob nor Merlin can speak to Alice), yet **OIP_+** is still powerful enough to simulate any streaming interactive proof, and **OIP** powerful enough to simulate all *known* streaming interactive proofs. Our study revealed that the online nature of these communication models leads them to behave very differently from classical interactive proofs, and allowed us to establish strong limitations on the power of existing techniques for developing constant-round SIPs. It also yielded a surprising characterization of the communication complexity class **AM** in terms of online communication models (namely, $\mathbf{AM} = \mathbf{OIP}^{[4]} = \mathbf{OIP}_+^{[4]}$). We believe this characterization may prove useful in establishing non-trivial **AM** lower bounds, a problem that has been identified [23] as an important "first step" toward resolving the $\mathbf{\Pi_2} \neq \mathbf{\Sigma_2}$ problem in two-party communication complexity, one of the most important problems left open by Babai et al. [5].

Many questions remain for future work, but here we highlight just one: proving a superlogarithmic lower bound on the $\mathbf{OIP}^{[2]}_+$ communication cost of an explicit function. Progress on this question would yield the first superlogarithmic lower bounds on the cost of two-round SIPs. Moreover, we have shown that standard techniques easily establish that $\mathbf{OIP}^{[2]}_+$ is a subset of $\mathbf{AM}$, but have been unable to prove any superlogarithmic lower bounds against $\mathbf{OIP}^{[2]}_+$ protocols. Proving $\mathbf{OIP}^{[2]}_+$ lower bounds therefore represents an important (and potentially tractable) "zeroth step" toward resolving $\mathbf{\Pi_2} \neq \mathbf{\Sigma_2}$.

References

1   Scott Aaronson. Qma/qpoly ⊆ pspace/poly: De-merlinizing quantum protocols. In *21st Annual IEEE Conference on Computational Complexity (CCC 2006), 16-20 July 2006, Prague, Czech Republic*, pages 261–273. IEEE Computer Society, 2006.

2   Scott Aaronson and Avi Wigderson. Algebrization: A new barrier in complexity theory. *TOCT*, 1(1), 2009.

3   Farid Ablayev. Lower bounds for one-way probabilistic communication complexity and their application to space complexity. *Theoretical Computer Science*, 175(2):139–159, 1996.

4   Noga Alon, Yossi Matias, and Mario Szegedy. The space complexity of approximating the frequency moments. *J. Comput. Syst. Sci.*, 58(1):137–147, 1999.

5   Laszlo Babai, Peter Frankl, and Janos Simon. Complexity classes in communication complexity theory. In *Foundations of Computer Science, 1986., 27th Annual Symposium on*, pages 337–347, Oct 1986.

6   László Babai and Shlomo Moran. Arthur-merlin games: A randomized proof system, and a hierarchy of complexity class. *J. Comput. Syst. Sci.*, 36(2):254–276, April 1988.

7   Dan Boneh, Giovanni Di Crescenzo, Rafail Ostrovsky, and Giuseppe Persiano. Public key encryption with keyword search. In *Advances in Cryptology – EUROCRYPT 2004, International Conference on the Theory and Applications of Cryptographic Techniques, Interlaken, Switzerland, May 2-6, 2004, Proceedings*, volume 3027 of *Lecture Notes in Computer Science*, pages 506–522. Springer, 2004.

8   Amit Chakrabarti, Graham Cormode, Navin Goyal, and Justin Thaler. Annotations for sparse data streams. In Chandra Chekuri, editor, *Proceedings of the Twenty-Fifth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2014, Portland, Oregon, USA, January 5-7, 2014*, pages 687–706. SIAM, 2014.

9   Amit Chakrabarti, Graham Cormode, Andrew McGregor, and Justin Thaler. Annotations in data streams. *ACM Transactions on Algorithms*, 11(1):7, 2014. A preliminary version of this paper by A. Chakrabarti, G. Cormode, and A. McGregor appeared in *ICALP* 2009.

10 Kai-Min Chung, Yael Tauman Kalai, Feng-Hao Liu, and Ran Raz. Memory delegation. In Phillip Rogaway, editor, *Advances in Cryptology – CRYPTO 2011 – 31st Annual Cryptology Conference, Santa Barbara, CA, USA, August 14-18, 2011. Proceedings*, volume 6841 of *Lecture Notes in Computer Science*, pages 151–168. Springer, 2011.

11 Raphaël Clifford, Klim Efremenko, Ely Porat, and Amir Rothschild. From coding theory to efficient pattern matching. In Claire Mathieu, editor, *Proceedings of the Twentieth*

Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2009, New York, NY, USA, January 4-6, 2009, pages 778–784. SIAM, 2009.

12   A. Condon. The complexity of space bounded interactive proof systems, 1993.

13   Graham Cormode, Michael Mitzenmacher, and Justin Thaler. Practical verified computation with streaming interactive proofs. In Shafi Goldwasser, editor, *Innovations in Theoretical Computer Science 2012, Cambridge, MA, USA, January 8-10, 2012*, pages 90–112. ACM, 2012.

14   Graham Cormode, Michael Mitzenmacher, and Justin Thaler. Streaming graph computations with a helpful advisor. *Algorithmica*, 65(2):409–442, 2013.

15   Graham Cormode, Justin Thaler, and Ke Yi. Verifying computations with streaming interactive proofs. *PVLDB*, 5(1):25–36, 2011.

16   S Goldwasser and M Sipser. Private coins versus public coins in interactive proof systems. In *Proceedings of the Eighteenth Annual ACM Symposium on Theory of Computing*, STOC'86, pages 59–68, New York, NY, USA, 1986. ACM.

17   Shafi Goldwasser, Dan Gutfreund, Alexander Healy, Tali Kaufman, and Guy N. Rothblum. A (de)constructive approach to program checking. In *Proceedings of the Fortieth Annual ACM Symposium on Theory of Computing*, STOC'08, pages 143–152, New York, NY, USA, 2008. ACM.

18   Shafi Goldwasser, Yael Tauman Kalai, and Guy N. Rothblum. Delegating computation: Interactive proofs for muggles. In *Proceedings of the Fortieth Annual ACM Symposium on Theory of Computing*, STOC'08, pages 113–122, New York, NY, USA, 2008. ACM.

19   Tom Gur and Ran Raz. Arthur-merlin streaming complexity. In FedorV. Fomin, Rusins Freivalds, Marta Kwiatkowska, and David Peleg, editors, *Automata, Languages, and Programming*, volume 7965 of *Lecture Notes in Computer Science*, pages 528–539. Springer Berlin Heidelberg, 2013.

20   Adam Kalai. Efficient pattern-matching with don't cares. In David Eppstein, editor, *Proceedings of the Thirteenth Annual ACM-SIAM Symposium on Discrete Algorithms, January 6-8, 2002, San Francisco, CA, USA.*, pages 655–656. ACM/SIAM, 2002.

21   Bala Kalyanasundaram and Georg Schintger. The probabilistic communication complexity of set intersection. *SIAM J. Discret. Math.*, 5(4):545–557, November 1992.

22   Hartmut Klauck. Rectangle size bounds and threshold covers in communication complexity. In *18th Annual IEEE Conference on Computational Complexity (Complexity 2003), 7-10 July 2003, Aarhus, Denmark*, pages 118–134. IEEE Computer Society, 2003.

23   Hartmut Klauck. On arthur merlin games in communication complexity. In *Proceedings of the 26th Annual IEEE Conference on Computational Complexity, CCC 2011, San Jose, California, June 8-10, 2011*, pages 189–199. IEEE Computer Society, 2011.

24   Hartmut Klauck and Ved Prakash. Streaming computations with a loquacious prover. In *Proceedings of the 4th Conference on Innovations in Theoretical Computer Science*, ITCS'13, pages 305–320, New York, NY, USA, 2013. ACM.

25   Ilan Kremer, Noam Nisan, and Dana Ron. On randomized one-round communication complexity. In *Proceedings of the Twenty-seventh Annual ACM Symposium on Theory of Computing*, STOC'95, pages 596–605, New York, NY, USA, 1995. ACM.

26   Satyanarayana V. Lokam. Spectral methods for matrix rigidity with applications to size–depth trade-offs and communication complexity. *Journal of Computer and System Sciences*, 63(3):449–473, 2001.

27   Carsten Lund, Lance Fortnow, Howard Karloff, and Noam Nisan. Algebraic methods for interactive proof systems. *J. ACM*, 39(4):859–868, October 1992.

28   Michael Mitzenmacher and Eli Upfal. *Probability and Computing: Randomized Algorithms and Probabilistic Analysis*. Cambridge University Press, New York, NY, USA, 2005.

29   Ilan Newman and Mario Szegedy. Public vs. private coin flips in one round communication games (extended abstract). In Gary L. Miller, editor, *Proceedings of the Twenty-Eighth Annual ACM Symposium on the Theory of Computing, Philadelphia, Pennsylvania, USA, May 22-24, 1996*, pages 561–570. ACM, 1996.

30   Noam Nisan and Avi Widgerson. Rounds in communication complexity revisited. In *Proceedings of the Twenty-third Annual ACM Symposium on Theory of Computing*, STOC'91, pages 419–429, New York, NY, USA, 1991. ACM.

31   Y. Ofman. On the algorithmic complexity of discrete functions. *Doklady Akademii Nauk SSSR*, 145:48–51, 1962. English Translation in Soviet Physics Doklady 7: 589-591, 1963.

32   Charalampos Papamanthou, Elaine Shi, Roberto Tamassia, and Ke Yi. Streaming authenticated data structures. In Thomas Johansson and Phong Q. Nguyen, editors, *Advances in Cryptology – EUROCRYPT 2013, 32nd Annual International Conference on the Theory and Applications of Cryptographic Techniques, Athens, Greece, May 26-30, 2013. Proceedings*, volume 7881 of *Lecture Notes in Computer Science*, pages 353–370. Springer, 2013.

33   Ran Raz. Quantum information and the PCP theorem. *Algorithmica*, 55(3):462–489, 2009.

34   Dominique Schröder and Heike Schröder. Verifiable data streaming. In Ting Yu, George Danezis, and Virgil D. Gligor, editors, *the ACM Conference on Computer and Communications Security, CCS'12, Raleigh, NC, USA, October 16-18, 2012*, pages 953–964. ACM, 2012.

35   Justin Thaler. Time-optimal interactive proofs for circuit evaluation. In Ran Canetti and Juan A. Garay, editors, *Advances in Cryptology – CRYPTO 2013 – 33rd Annual Cryptology Conference, Santa Barbara, CA, USA, August 18-22, 2013. Proceedings, Part II*, volume 8043 of *Lecture Notes in Computer Science*, pages 71–89. Springer, 2013.

36   Victor Vu, Srinath T. V. Setty, Andrew J. Blumberg, and Michael Walfish. A hybrid architecture for interactive verifiable computation. In *2013 IEEE Symposium on Security and Privacy, SP 2013, Berkeley, CA, USA, May 19-22, 2013*, pages 223–237. IEEE Computer Society, 2013.

37   C.S. Wallace. A suggestion for a fast multiplier. *Electronic Computers, IEEE Transactions on*, EC-13(1):14–17, Feb 1964.