On Interactive Proofs with a Laconic Prover

(Extended Abstract)

Oded Goldreich[1,⋆], Salil Vadhan[2,⋆⋆], and Avi Wigderson[3,⋆⋆⋆]

[1] Department of Computer Science, Weizmann Institute of Science, Rehovot, ISRAEL
`oded@wisdom.weizmann.ac.il`

[2] Division of Engineering & Applied Sciences, Harvard University, Cambridge, MA
`salil@eecs.harvard.edu, http://eecs.harvard.edu/~salil`

[3] School of Mathematics, Institute for Advanced Study, Princeton, NJ
`avi@ias.edu`

Abstract. We continue the investigation of interactive proofs with bounded communication, as initiated by Goldreich and Håstad (IPL 1998). Let L be a language that has an interactive proof in which the prover sends few (say b) bits to the verifier. We prove that the complement $\bar{L}$ has a *constant-round* interactive proof of complexity that depends only exponentially on b. This provides the first evidence that for **NP**-complete languages, we cannot expect interactive provers to be much more "laconic" than the standard **NP** proof.
When the proof system is further restricted (*e.g.,* when $b = 1$, or when we have perfect completeness), we get significantly better upper bounds on the complexity of $\bar{L}$.

Keywords: interactive proofs, Arthur-Merlin games, sampling protocols, statistical zero knowledge, game theory

1 Introduction

Interactive proof systems were introduce by Goldwasser, Micali and Rackoff [GMR89] in order to capture the most general way in which one party can *efficiently verify* claims made by another, more powerful party.[1] That is, interactive proof systems are two-party randomized protocols through which a computationally

---

⋆ Supported by the MINERVA Foundation.

⋆⋆ Work done while at the Institute for Advanced Study, Princeton, NJ, supported by an NSF Mathematical Sciences Postdoctoral Research Fellowship.

⋆⋆⋆ Partially supported by NSF grants CCR-9987845 and CCR-9987077.

[1] Arthur-Merlin games, introduced by Babai [Bab85], are a special type on interactive proofs in which the verifier is restricted to send the outcome of each coin it tosses. Such proof systems are also called *public coin*, and are known to be as expressive as general interactive proofs [GS89]. We warn that the latter assertion refers to the entire class but not to refined complexity measures such as the number of bits sent by the prover (considered below).

unbounded prover can convince a probabilistic polynomial-time verifier of the membership of a common input in a predetermined language. Thus, interactive proof systems generalize and contain as a special case the traditional "**NP**-proof systems" (in which verification is deterministic and "non-interactive").

It is well-known that this generalization buys us a lot: The *IP Characterization Theorem* of Lund, Fortnow, Karloff, Nisan and Shamir [LFKN92, Sha92] states that every language in **PSPACE** has an interactive proof system, and it is easy to see that only languages in **PSPACE** have interactive proof systems.

It is well-known that the strong expressive power of interactive proofs is largely due to the presence of interaction. In particular, interactive proofs in which a single message is sent (like in **NP**-proofs) yield a complexity class (known as **MA**) that seems very close to **NP**. It is interesting to explore what happens between these extremes of unbounded interaction and no interaction. That is, *what is the expressive power of interactive proofs that utilize a bounded, but nonzero, amount of interaction?*

Interactive Proofs with Few Messages. The earliest investigations of the above question examined the *message complexity* of interactive proofs, *i.e.,* the number of messages exchanged. (Sometimes, we refer to *rounds*, which are a pair of verifier-prover messages.) The Speedup Theorem of Babai and Moran [BM88] (together with [GS89]) shows that the number of messages in an interactive proof can be always be reduced by a constant factor (provided the number of messages remains at least 2). On the other hand, there is a large gap between constant-round interactive proofs and unrestricted interactive proofs. As mentioned above, all of **PSPACE** has a general interactive proof [LFKN92, Sha92]. In contrast, the class **AM** of problems with constant-round interactive proofs is viewed as being relatively close to **NP**. Specifically, **AM** lies in the second level of the polynomial-time hierarchy [BM88], cannot contain **coNP** unless the polynomial-time hierarchy collapses [BHZ87], and actually equals **NP** under plausible circuit complexity assumptions [AK97, KvM99, MV99].

Laconic Provers. A more refined investigation of the above question was initiated by Goldreich and Håstad [GH98], who gave bounds on the complexity of languages possessing interactive proofs with various restrictions on the *number of bits* of communication and/or randomness used. One of the restrictions they considered, and the main focus of our investigation, limits the number of bits sent from the prover to the verifier by some bound b. That is, what languages can be proven by "laconic" provers?

Since the prover is trying to convey something to the verifier, this seems to be the most interesting direction of communication. Moreover, for applications of interactive proofs (*e.g.,* in cryptographic protocols), it models the common situation in which communication is more expensive in one direction (*e.g.,* if the prover is a handheld wireless device).

On one hand, we know of interactive proofs for several "hard" problems (Quadratic Nonresiduosity [GMR89], Graph Nonisomorphism [GMW91], and others [GK93, GG00, SV97]) in which the communication from the prover to the verifier is severely bounded (in fact, to one bit). On the other hand, no such proof systems were known for **NP**-complete problems, nor was there any indication of impossibility (except when additional constraints are imposed [GH98]). In this work, we provide strong evidence of impossibility.

Our Results. Consider interactive proofs in which the prover sends at most $b = b(n)$ bits to the verifier on inputs of length n. Goldreich and Håstad [GH98, Thm. 4] placed such languages in $\mathbf{BPTIME}^{\mathbf{NP}}(T)$, where $T = \mathrm{poly}(n) \cdot 2^{\mathrm{poly}(b)}$, which clearly implies nothing for languages in **NP**. In contrast, we show that the *complements* of such languages have *constant-round* interactive proofs of complexity T (*i.e.,* the verifier's computation time and the total communication is bounded by T). In particular, **NP**-complete problems cannot have interactive proofs in which in which the prover sends at most polylogarithmically many bits to the verifier unless **coNP** is in the quasipolynomial analogue of **AM**. In fact, assuming **NP** has constant-round interactive proofs with logarithmic prover-to-verifier communication we conclude $\mathbf{coNP} \subset \mathbf{AM}$. As mentioned above, this is highly unlikely.

We obtain stronger results in two special cases:

1. We show that if a language has an interactive proof of perfect completeness (*i.e.,* , zero error probability on YES instances) in which the prover sends at most $b(n)$ bits, then it is in $\mathbf{coNTIME}(T)$, where $T(n) = 2^{b(n)} \cdot \mathrm{poly}(n)$. Thus, unless $\mathbf{NP} = \mathbf{coNP}$, NP-complete languages cannot have interactive proof systems of perfect completeness in which the prover sends at most logarithmically many bits.
2. We show that if a language has an interactive proof in which the prover sends a single bit (with some restrictions on the error probabilities), then it has a statistical zero-knowledge interactive proof; that is, is in the class **SZK**. This is a stronger conclusion than our main result because $\mathbf{SZK} \subseteq \mathbf{AM} \cap \mathbf{coAM}$, as shown by Fortnow [For89] and Aiello and Håstad [AH91]. Recalling that Sahai and Vadhan [SV97] showed that any language in **SZK** has an interactive proof in which the prover sends a single bit, we obtain a surprising equivalence between these two classes.[2]

Lastly, we mention one easy, but apparently new, observation regarding message complexity. A question that is left open by the results mentioned earlier is what happens "in between" constant rounds and polynomially many rounds. Phrased differently, can the Speedup Theorem of Babai and Moran be improved to show that $m(n)$-message interactive proofs are no more powerful than $m'(n)$-message ones for some $m' = o(m)$? By combining careful parameterizations of [LFKN92, BM88], we observe that such an improvement is unlikely. More precisely, for every nice function m, we show that there is a language which has an

[2] In addition, if the error probabilities are sufficiently small, we also are able to reduce interactive proofs in which the prover sends a single *message* of several bits (*e.g.,* $O(\log\log n)$ bits) to the 1-bit case above. But we omit these results from this extended abstract due to space constraints.

$m(n)$-message interactive proof but not an $o(m(n))$-message one, provided that #SAT is not contained in the subexponential analogue of **coAM**.

Additional Related Work. It should be noted that the results of Goldreich and Håstad are significantly stronger when further restrictions are imposed in addition to making the prover laconic. In particular, they obtain an upper bound of $\mathbf{BPTIME}(T)$ (rather than $\mathbf{BPTIME}^{\mathbf{NP}}(T)$), with $T = 2^{\mathrm{poly}(b)} \cdot \mathrm{poly}(n)$ for languages possessing either of the following kinds of interactive proofs: (a) *public-coin* proofs in which the prover sends at most b bits, (b) proofs in which the communication *in both directions* is bounded by b.

There has also been a body of research on the expressive power of *multi-prover interactive proofs* (MIP's) and *probabilistically checkable proofs* (PCP's) with low communication, because of the importance of the communication parameter in their applications to inapproximability. In particular, Bellare, Goldreich, and Sudan [BGS98] give negative results about the expressive power of "laconic" PCP's and MIP's. One-query probabilistically checkable proofs are equivalent to interactive proofs in which the prover sends a single message, so our results provide bounds on the former.

Our work is also related to work on *knowledge complexity.* Knowledge complexity, proposed by [GMR89], aims to measure how much "knowledge" is leaked from the prover to the verifier in an interactive proof. Several measures of knowledge complexity were proposed by Goldreich and Petrank [GP99], and series of works provided upper bounds on the complexity of languages having interactive proofs with low knowledge complexity [GP99, GOP98, PT96, SV97]. These results are related to, but incomparable to ours.

For example, Petrank and Tardos [PT96] showed that languages having knowledge complexity $k = O(\log n)$ are contained in $\mathbf{AM} \cap \mathbf{coAM}$. While it is true that the knowledge complexity of an interactive proof is bounded by the amount of prover-to-verifier communication, their result does not yield anything interesting for laconic interactive proofs. The reason is that their result only applies to interactive proofs with error probabilities significantly smaller than 2^{-k}, and it is easy to see that interactive proofs with prover-to-verifier communication $k = O(\log n)$ error probability $\ll 2^{-k}$ only capture **BPP** (and hence are uninteresting). Our results apply even for constant error probabilities.

Sahai and Vadhan [SV97] (improving [GP99]) showed that languages with logarithmic knowledge complexity in the "hint sense" collapse to **SZK**, and their result applies even if the error probabilities are constant. However, this is also incomparable to ours, for the "hint sense" is the one measure of knowledge complexity which is *not* bounded by the prover-to-verifier communication.

Finally, it is important to note that the situation is dramatically different for *argument systems* [BCC88] (also known as *computationally sound proofs*). These are like interactive proofs, but the soundness condition is restricted to polynomial-time provers. Kilian [Kil92] showed that **NP** has laconic argument systems if strong collision-resistant hash functions exist. Specifically, under a strong enough (but still plausible) assumption, **NP** has *public-coin* arguments in which the verifier's randomness and the communication in both directions is *polylogarithmic*. Combined with [GH98], this provides a strong separation between the efficiency of arguments versus interactive proofs for **NP**; and our results extend this separation to the case that only the prover-to-verifier communication is counted (and the interactive proof is not required to be public coin).

2 Preliminaries

We assume that the reader is familiar with the basic concepts underlying interactive proofs (and public-coin interactive proofs) (see *e.g.,* [Sip97, Gol99, Vad00]). Throughout, we work with interactive proofs for *promise problems* rather than languages. More precisely, a promise problem $\Pi = (\Pi_Y, \Pi_N)$ is a pair of disjoint sets of strings, corresponding to YES and NO instances, respectively. In other words, a promise problem is simply a decision problem in which some inputs are excluded. The definition of interactive proofs is extended to promise problems in the natural way: we require that when the input is a YES instance, the prover convinces the verifier to accept with high probability (*completeness*); and when the input is a NO instance, the verifier accepts with low probability no matter what strategy the prover follows (*soundness*). Working with promise problems rather than languages only makes our results stronger (except for one direction of Theorem 4.4).

We denote by $\mathbf{IP}(b,m)$ (resp., $\mathbf{AM}(b,m)$) the class of problems having interactive proofs (resp., public-coin interactive proofs) in which the prover sends a total of at most b bits, and the total number of messages exchanged (in both directions) is at most m. Note that b and m are integer functions of the common input length, denoted n. When b is not polynomial in n, it will be understood that we talk of a generalization in which the verifier is allowed time polynomial in b and n (rather than just in n). Unless specified differently, we refer to proof systems with completeness probability 2/3 and soundness probability 1/3.

We denote $\mathbf{IP}(b) = \mathbf{IP}(b, 2b)$; that is, making only the trivial bound on the number of messages exchanged. We denote by $\mathbf{IP}^+$ the analogue of $\mathbf{IP}$ when the proof system has perfect completeness (*i.e.,* completeness probability 1). The class of problems with constant-round interactive proofs is denoted $\mathbf{AM} \stackrel{\text{def}}{=} \mathbf{AM}(\text{poly}(n), 2) = \mathbf{IP}(\text{poly}(n), O(1))$. (The second equality is by Thms 2.3 and 2.4 below.) When we wish to specify the completeness probability $c = c(n)$ and soundness probability $s = s(n)$ we will use subscripts: $\mathbf{IP}_{c,s}$ and $\mathbf{AM}_{c,s}$.

Using the above notations, we recall the main results of Goldreich and Håstad, which are the starting point for our work.

Theorem 2.1 ([GH98]). $\mathbf{AM}(b,m) \subseteq \mathbf{BPTIME}(\text{poly}(2^b, m^m, n))$

Theorem 2.2 ([GH98]). $\mathbf{IP}(b,m) \subseteq \mathbf{BPTIME}(\text{poly}(2^b, m^m, n))^{\mathbf{NP}}$

We also state some standard results that we will use:

Theorem 2.3 ([BM88]). $\mathbf{AM}(b,m) \subseteq \mathbf{AM}(b^2 \cdot \text{poly}(m), \lceil m/2 \rceil) \subseteq \mathbf{AM}((b \cdot m)^{O(m)}, 2)$.

Theorem 2.4 **([GS89]).** $\mathbf{IP}(b, m) \subseteq \mathbf{AM}(\mathrm{poly}(b, n), m)$.

Theorem 2.5 **([BHZ87]).** *If* $\mathbf{coNP} \subseteq \mathbf{AM}(b, 2)$*, then* $\mathbf{\Sigma_2} \subseteq \mathbf{\Pi_2}(\mathrm{poly}(n, b))$. *In particular, if* $\mathbf{coNP} \subset \mathbf{AM}$*, then the polynomial-time hierarchy collapses to* $\mathbf{PH} = \mathbf{\Sigma_2} = \mathbf{\Pi_2}$.

Above and throughout the paper, $\mathbf{\Sigma_i}(t(n))$ (resp., $\mathbf{\Pi_i}(t(n))$) denotes the class of problems accepted by $t(n)$-time alternating Turing machines with i alternations beginning with an existential (resp., universal) quantifier. Thus, $\mathbf{\Sigma_i} \stackrel{\mathrm{def}}{=} \mathbf{\Sigma_i}(\mathrm{poly}(n))$ and $\mathbf{\Pi_i} \stackrel{\mathrm{def}}{=} \mathbf{\Pi_i}(\mathrm{poly}(n))$ comprise the i'th level of the polynomial-time hierarchy.

We will also consider **SZK**, the class of problems possessing statistical zero-knowledge interactive proofs. Rather than review the definition here, we will instead use a recent characterization of it in terms of complete problems which will suffice for our purposes. For distributions X and Y, let $\Delta(X, Y)$ denote their *statistical difference* (or *variation distance, i.e.,* $\Delta(X, Y) = \max_S |\Pr[X \in S] - \Pr[Y \in S]|$. We will consider distributions specified by circuits which sample from them. More precisely, a circuit with m input gates and n output gates can be viewed as a sampling algorithm for the distribution on $\{0,1\}^n$ induced by evaluating the circuit on m random input bits. Statistical Difference is the promise problem $\mathrm{SD} = (\mathrm{SD}_Y, \mathrm{SD}_N)$, where

$$\begin{aligned} \mathrm{SD}_Y &= \{(X, Y) \ : \ \Delta(X, Y) \geq 2/3\} \\ \mathrm{SD}_N &= \{(X, Y) \ : \ \Delta(X, Y) \leq 1/3\}, \end{aligned}$$

where X and Y are probability distributions specified by circuits which sample from them. More generally, for any $1 \geq \alpha > \beta \geq 0$, we will consider variants $\mathrm{SD}^{\alpha,\beta}$, where the thresholds of $2/3$ and $1/3$ are replaced with α and β respectively.

Theorem 2.6 **([SV97]).** *For any constants* $1 > \alpha^2 > \beta > 0$, $\mathrm{SD}^{\alpha,\beta}$ *is complete for* **SZK**.

The following results about **SZK** are also relevant to us.

Theorem 2.7 **([For89, AH91]).** $\mathbf{SZK} \subseteq \mathbf{AM} \cap \mathbf{coAM}$.

Theorem 2.8 **([Oka00]).** **SZK** *is closed under complement.*

Theorem 2.9 **([SV97]).** $\mathbf{SZK} \subseteq \mathbf{IP}_{1-2^{-n},1/2}(1)$.

3 Formal Statement of Results

We improve over Theorem 2.2, and address most of the open problems suggested in [GH98, Sec. 3]. Our main results are listed below.

For one bit of prover-to-verifier communication, we obtain a collapse to **SZK**.

Theorem 3.1. *For every pair of constants c, s such that $1 > c^2 > s > c/2 > 0$, $\mathbf{IP}_{c,s}(1) = \mathbf{SZK}$.*

With Theorem 2.8, this gives:

Corollary 3.2. *For every c, s as in Thm. 3.1, $\mathbf{IP}_{c,s}(1)$ is closed under complement.*

For more rounds of communication, we first obtain the following result for interactive proofs with perfect completeness (denoted by $\mathbf{IP}^+$):

Theorem 3.3. $\mathbf{IP}^+(b) \subseteq \mathbf{coNTIME}(2^b \cdot \mathrm{poly}(n))$. *In particular,* $\mathbf{IP}^+(O(\log n)) \subseteq \mathbf{coNP}$.

In the general case (*i.e.,* with imperfect completeness), we prove:

Theorem 3.4. $\mathbf{IP}(b, m) \subseteq \mathbf{coAM}(2^b \cdot \mathrm{poly}(m^m, n), O(m))$. *In particular,* $\mathbf{IP}(O(\log n), m) \subseteq \mathbf{coAM}(\mathrm{poly}(n), O(m))$, *for* $m = O(\log n / \log\log n)$,

The above theorems provide first evidence that $\mathbf{NP}$-complete problems cannot have interactive proof systems in which the prover sends very few bits. Further evidence toward this claim is obtained by applying Theorems 2.3 and 2.5:

Corollary 3.5. $\mathbf{IP}(b, m) \subseteq \mathbf{coAM}(\mathrm{poly}(2^b, m^m, n)^m, 2)$. *In particular,* $\mathbf{IP}(O(\log n), O(1)) \subset \mathbf{coAM}$ *and* $\mathbf{IP}(\mathrm{polylog}\, n) \subset \widetilde{\mathbf{coAM}}$.

Corollary 3.6. $\mathbf{NP} \not\subseteq \mathbf{IP}(O(\log n), O(1))$ *unless the polynomial-time hierarchy collapses (to $\mathbf{\Sigma_2} = \mathbf{\Pi_2}$).* $\mathbf{NP} \not\subseteq \mathbf{IP}(\mathrm{polylog}\, n)$ *unless* $\mathbf{\Sigma_2} \subseteq \widetilde{\mathbf{\Pi}}_2$.

Above, $\widetilde{\mathbf{coAM}}$ and $\widetilde{\mathbf{\Pi}}_2$ denote the quasipolynomial-time ($2^{\mathrm{polylog}\, n}$) analogues of $\mathbf{coAM}$ and $\mathbf{\Pi_2}$.

Finally, we state our result on message complexity.

Theorem 3.7. *Let $m(n) \leq n/\log n$ be any "nice" growing function. Then* $\mathbf{AM}(\mathrm{poly}(n), m(n)) \neq \mathbf{AM}(\mathrm{poly}(n), o(m(n)))$ *unless* $\#\mathrm{SAT} \in \mathbf{AM}(2^{o(n)}, 2)$.

Note that, by Theorem 2.4, it is irrelevant whether we use $\mathbf{IP}$ or $\mathbf{AM}$ in this theorem.

Due to space constraints, we only present proofs of Theorems 3.1 and 3.3 in this extended abstract. The proof of our main result (Theorem 3.4) is significantly more involved, and will be given in the full version of the paper.

4 Extremely Laconic Provers (Saying Only One Bit)

In this section, we prove Theorem 3.1. The proof is based on the following lemma, along with previous results.

Lemma 4.1. *Every problem in $\mathbf{IP}_{c,s}(1)$ reduces to $\mathrm{SD}^{c,s}$.*

Proof. Let (P, V) be an interactive proof for some problem so that the prover sends a single bit during the entire interaction. We may thus assume that on input x and internal coin tosses r, the verifier first sends a message $y = V_x(r)$, the prover answers with a bit $\sigma \in \{0, 1\}$, and the verifier decides whether to accept or reject by evaluating the predicate $V_x(r, \sigma) \in \{0, 1\}$.

A special case — unique answers. To demonstrate the main idea, we consider first the natural case in which for every pair (x, r) there exists *exactly one* σ such that $V_x(r, \sigma) = 1$. (Note that otherwise, the interaction on input x and verifier's internal coin tosses r is redundant, since the verifier's final decision is unaffected by it.) For this special case (which we refer to as *unique answers*), we will prove the following:

Claim 4.2. *If a problem has an* $\mathbf{IP}_{c,s}(1)$ *proof system with unique answers, then it reduces to* $\mathrm{SD}^{2c-1,2s-1}$.

Let $\sigma_x(r)$ denote the unique σ satisfying $V_x(r, \sigma) = 1$. The prover's ability to convince the verifier is related to the amount of information regarding $\sigma_x(r)$ that is revealed by $V_x(r)$. For example, if for some x, $\sigma_x(r)$ is determined by $V_x(r)$ then the prover can convince the verifier to accept x with probability 1 (by replying with $\sigma_x(r)$). If, on the other hand, for some x, $\sigma_x(r)$ is statistically independent of $V_x(r)$ (and unbiased), then there is no way for the prover to convince the verifier to accept x with probability higher than $1/2$. This suggests the reduction $x \mapsto (C_x^1, C_x^2)$, where $C_x^1(r) \stackrel{\text{def}}{=} (V_x(r), \sigma_x(r))$ and $C_x^2(r) \stackrel{\text{def}}{=} (V_x(r), \overline{\sigma_x}(r))$, where $\overline{b}$ denotes the complement of a bit b.

Now we relate the statistical difference between the distributions sampled by C_x^1 and C_x^2 to the maximum acceptance probability of the verifier. Since the first components of C_x^1 and C_x^2 are distributed identically, their statistical difference is exactly the average over the first component $V_x(r)$ of the statistical difference between the second components conditioned on $V_x(r)$. That is,

$$\Delta(C_x^1, C_x^2) = \mathop{\mathrm{E}}_{y \leftarrow V_x} \left[\Delta\left(\sigma_x|_y, \overline{\sigma_x}|_y\right)\right],$$

where $\sigma_x|_y$ denotes the distribution of $\sigma_x(r)$ when r is uniformly distributed among $\{r' : V_x(r') = y\}$. For any y and $b \in \{0, 1\}$, let $q_{b|y}$ denote the probability that $\sigma_x|_y = b$. Then, for any fixed y, $\Delta\left(\sigma_x|_y, \overline{\sigma_x}|_y\right) = |q_{1|y} - q_{0|y}| = 2q_y - 1$, where $q_y \stackrel{\text{def}}{=} \max_{b \in \{0,1\}}\{q_{b|y}\} \geq \frac{1}{2}$. So, we have:

$$\Delta(C_x^1, C_x^2) = \mathop{\mathrm{E}}_{y \leftarrow V_x} \left[2q_y - 1\right].$$

On the other hand, the optimal prover strategy in (P, V) is: upon receiving y, respond with b that maximizes $q_{b|y}$. When the prover follows this strategy, we have

$$\Pr[V \text{ accepts } x] = \mathop{\mathrm{E}}_{y \leftarrow V_x} \left[q_y\right].$$

Putting the last two equations together, we conclude that $\Delta(C_x^1, C_x^2) = 2 \cdot \Pr[V \text{ accepts } x] - 1$.[3] Thus if the proof system has completeness and soundness error bounds c and s, respectively, then the reduction maps instances to

[3] Note that under the hypothesis of the special case, for every x the prover may convince the verifier to accept x with probability at least $1/2$ (and so such a non-trivial proof system must have soundness at least $1/2$).

pairs having distance bounds $2c-1$ and $2s-1$, respectively.[4] This establishes Claim 4.2.

The general case. We now proceed to deal with the general case in which there may exist pairs (x,r) so that either both σ's or none of them satisfy $V_x(r,\sigma)=1$. We do so by reducing this general case to the special case.

Claim 4.3. *If a problem is in* $\mathbf{IP}_{c,s}(1)$*, then it has an* $\mathbf{IP}_{(1+c)/2,(1+s)/2}(1)$ *proof system with unique answers.*

Clearly, the lemma follows from this claim and the previous one, so we proceed to prove the claim.

> *Proof of claim.* Let (P,V) be a general $\mathbf{IP}_{c,s}$ proof system. Consider the following modified verifier strategy.
>
> $V'(x)$**:** Generate coin tosses r for the original verifier and do one of the following based on the number j of possible prover responses σ for which $V_x(r,\sigma)=1$.
>
> $[j=2]$ Send the prover a special message "`respond with 1`" and accept if the prover responds with 1.
>
> $[j=1]$ Randomly do one of the following (each with prob. 1/2):
> - Send the prover $y=V_x(r)$ and accept if the prover responds with the unique σ such that $V_x(r,\sigma)=1$.
> - Send the prover a special message "`respond with 1`" and accept if the prover responds with 1.
>
> $[j=0]$ Choose a random bit σ. Send the prover a special message "`guess my bit`" and accept if the prover responds with σ.
>
> Clearly, V' has unique answers. It can be shown that if an optimal prover makes V accept with probability δ, then an optimal prover makes V' accept with probability $(1+\delta)/2$. Claim 4.3 follows. □

■

Theorem 3.1 follows from Lemma 4.1, Theorem 2.6, and Theorem 2.9. Details will be given in the full version of the paper. The $c^2>s$ constraint in Theorem 3.1 is due to the analogous constraint in Theorem 2.6. Indeed, we can establish the following equivalence (also to be proven the full version of the paper):

Theorem 4.4. *The following are equivalent.*

1. *For every α,β such that $1>\alpha>\beta>0$,* $\mathrm{SD}^{\alpha,\beta}$ *is in* **SZK** *(and is therefore also complete).*
2. *For every c,s such that $1>c>s>c/2>0$,* $\mathbf{IP}_{c,s}(1)=\mathbf{SZK}$.

Finally, we remark that the condition $s>c/2$ in Theorems 3.1 and 4.4 is necessary, for $\mathbf{IP}_{c,s}(1)=\mathbf{BPP}$ for any $s<c/2$.

[4] Note that this relationship is reversed by the natural $\mathbf{IP}(1)$ system for $\mathrm{SD}^{\alpha,\beta}$ in which the verifier selects at random a single sample from one of the two distributions and asks the prover to guess which of the distributions this sample came from. If the distributions are at distance δ then the prover succeeds with probability $\frac{1}{2}+\frac{\delta}{2}$. Thus applying this proof system to $\mathrm{SD}^{2c-1,2s-1}$ we obtain completeness and soundness bounds c and s, respectively.

5 Laconic Provers with Perfect Completeness

In this section, we prove Theorem 3.3.

Theorem 3.3 (restated): If a problem Π has an interactive proof system with perfect completeness in which the prover-to-verifier communication is at most $b(\cdot)$ bits then $\Pi \in \mathbf{coNTIME}(2^{b(n)} \cdot \mathrm{poly}(n))$.

Proof. We take a slightly unusual look at the interactive proof system for Π, viewing it as a "progressively finite game" between two players P^* and V^*. P^* corresponds to the usual prover strategy and its aim is to make the original verifier accept the common input. V^* is a "cheating verifier" and its aim is to produce an interaction that looks legal and still makes the original verifier reject the common input.

To make this precise, let $b = b(n)$ be the bound on the prover-to-verifier communication in (P, V) on inputs of length n, and let $m = m(n)$ be the number of messages exchanged. Without loss of generality, we may assume that the V sends all its coin tosses in the last message. A *transcript* is a sequence of m strings, corresponding to (possible) messages exchanged between P and V. We call a transcript t *consistent* (for x) if every verifier message in t is the message V would have sent given input x, the previous messages in t, and the coin tosses specified by the last message in t. We call a consistent t *rejecting* if V would reject at the end of such an interaction.

Now, the game between P_x^* and V_x^* has the same structure as the interaction between P and V on input x: a total of m messages are exchanged and P_x^* is allowed to send at most b bits. The game between P_x^* and V_x^* yields a transcript t. We say that V_x^* *wins* if t is consistent and rejecting, and that P_x^* wins otherwise. We stress that V_x^* need not emulate the original verifier nor is it necessarily implemented in probabilistic polynomial time.

This constitutes a "perfect information finite game in extensive form" (also known as a "progressively finite game") and Zermelo's Theorem (*cf.,* [Tuc95, Sec 10.2]) says that exactly one of the two players has a *winning strategy* — that is, a (deterministic) strategy that will guarantee its victory no matter how the other party plays.

Using the perfect completeness condition, we infer that if the common input x is a YES instance then there exists a winning strategy for P_x^*. (This is because the optimal prover for the original interactive proof wins whenever V_x^* plays in a manner consistent with some sequence of coin tosses for the original verifier, and it wins by definition if the V_x^* plays inconsistently with any such sequence.) On the other hand, by the soundness condition, if the common input is a NO instance then there exists no winning strategy for P_x^*. (This is because in this case no prover strategy can convince the original verifier with probability 1.) By the above, it follows that whenever the common input is a NO instance there exists a winning strategy for V_x^*.

Thus, a proof that x is a NO instance consists of a winning strategy for V_x^*. Such strategy is a function mapping partial transcripts of P_x^* messages to the next V_x^* message. Thus, such a strategy is fully specified by a function from $\cup_{i=0}^{b}\{0,1\}^i$ to $\{0,1\}^{\mathrm{poly}(n)}$, and has description length $\mathrm{poly}(n)\cdot 2^{b(n)+1}$. To verify that such a function constitutes a winning strategy for V_x^*, one merely tries all possible deterministic strategies for the P_x^* (i.e., all possible $b(n)$-bit long strings). The theorem follows. ■

References

[AH91] William Aiello and Johan Håstad. Statistical zero-knowledge languages can be recognized in two rounds. *Journal of Computer and System Sciences*, 42(3):327–345, June 1991.

[AK97] V. Arvind and J. Köbler. On resource-bounded measure and pseudorandomness. In *Proceedings of the 17th Conference on Foundations of Software Technology and Theoretical Computer Science*, pages 235–249. LNCS 1346, Springer-Verlag, 1997.

[Bab85] László Babai. Trading group theory for randomness. In *Proceedings of the Seventeenth Annual ACM Symposium on Theory of Computing*, pages 421–429, Providence, Rhode Island, 6–8 May 1985.

[BM88] László Babai and Shlomo Moran. Arthur-Merlin games: A randomized proof system and a hierarchy of complexity classes. *Journal of Computer and System Sciences*, 36:254–276, 1988.

[BGS98] Mihir Bellare, Oded Goldreich, and Madhu Sudan. Free bits, PCPs, and nonapproximability—towards tight results. *SIAM Journal on Computing*, 27(3):804–915 (electronic), 1998.

[BHZ87] Ravi B. Boppana, Johan Håstad, and Stathis Zachos. Does co-NP have short interactive proofs? *Information Processing Letters*, 25:127–132, 1987.

[BCC88] Gilles Brassard, David Chaum, and Claude Crépeau. Minimum disclosure proofs of knowledge. *Journal of Computer and System Sciences*, 37(2):156–189, October 1988.

[For89] Lance Fortnow. The complexity of perfect zero-knowledge. In Silvio Micali, editor, *Advances in Computing Research*, volume 5, pages 327–343. JAC Press, Inc., 1989.

[Gol99] Oded Goldreich. *Modern Cryptography, Probabilistic Proofs, and Pseudorandomness*. Number 17 in Algorithms and Combinatorics. Springer-Verlag, 1999.

[GG00] Oded Goldreich and Shafi Goldwasser. On the limits of nonapproximability of lattice problems. *Journal of Computer and System Sciences*, 60(3):540–563, 2000.

[GH98] Oded Goldreich and Johan Håstad. On the complexity of interactive proofs with bounded communication. *Information Processing Letters*, 67(4):205–214, 1998.

[GK93] Oded Goldreich and Eyal Kushilevitz. A perfect zero-knowledge proof system for a problem equivalent to the discrete logarithm. *Journal of Cryptology*, 6:97–116, 1993.

[GMW91] Oded Goldreich, Silvio Micali, and Avi Wigderson. Proofs that yield nothing but their validity or all languages in NP have zero-knowledge proof systems. *Journal of the ACM*, 38(1):691–729, 1991.

[GOP98] Oded Goldreich, Rafail Ostrovsky, and Erez Petrank. Computational complexity and knowledge complexity. *SIAM Journal on Computing*, 27(4):1116–1141, August 1998.

[GP99] Oded Goldreich and Erez Petrank. Quantifying knowledge complexity. *Computational Complexity*, 8(1):50–98, 1999.

[GMR89] Shafi Goldwasser, Silvio Micali, and Charles Rackoff. The knowledge complexity of interactive proof systems. *SIAM Journal on Computing*, 18(1):186–208, February 1989.

[GS89] Shafi Goldwasser and Michael Sipser. Private coins versus public coins in interactive proof systems. In Silvio Micali, editor, *Advances in Computing Research*, volume 5, pages 73–90. JAC Press, Inc., 1989.

[Kil92] Joe Kilian. A note on efficient zero-knowledge proofs and arguments (extended abstract). In *Proceedings of the Twenty-Fourth Annual ACM Symposium on the Theory of Computing*, pages 723–732, Victoria, British Columbia, Canada, 4–6 May 1992.

[KvM99] Adam R. Klivans and Dieter van Melkebeek. Graph nonisomorphism has subexponential size proofs unless the polynomial-time hierarchy collapses. In *Proceedings of the Thirty-first Annual ACM Symposium on Theory of Computing*, pages 659–667, Atlanta, 1–4 May 1999.

[LFKN92] Carsten Lund, Lance Fortnow, Howard Karloff, and Noam Nisan. Algebraic methods for interactive proof systems. *Journal of the ACM*, 39(4):859–868, October 1992.

[MV99] Peter Bro Miltersen and N.V. Vinodchandran. Derandomizing Arthur–Merlin games using hitting sets. In *40th Annual Symposium on Foundations of Computer Science*, New York, NY, 17–19 October 1999. IEEE.

[Oka00] Tatsuaki Okamoto. On relationships between statistical zero-knowledge proofs. *Journal of Computer and System Sciences*, 60(1):47–108, February 2000.

[PT96] Erez Petrank and Gábor Tardos. On the knowledge complexity of $\mathcal{NP}$. In *37th Annual Symposium on Foundations of Computer Science*, pages 494–503, Burlington, Vermont, 14–16 October 1996. IEEE.

[SV97] Amit Sahai and Salil P. Vadhan. A complete promise problem for statistical zero-knowledge. In *38th Annual Symposium on Foundations of Computer Science*, pages 448–457, Miami Beach, Florida, 20–22 October 1997. IEEE.

[Sha92] Adi Shamir. IP = PSPACE. *Journal of the ACM*, 39(4):869–877, October 1992.

[Sip97] Michael Sipser. *Introduction to the Theory of Computation*. PWS Publishing, 1997.

[Tuc95] Alan Tucker. *Applied combinatorics*. John Wiley & Sons Inc., New York, third edition, 1995.

[Vad00] Salil Vadhan. Probabilistic proof systems I: Interactive and zero-knowledge proofs. Lecture Notes from the *IAS/PCMI Graduate Summer School on Computational Complexity*, August 2000. Available from `http://eecs.harvard.edu/~salil/`.