

Advisor-Verifier-Prover Protocols Are Lengthy

One hypothesis that would give super-polynomial lower bounds for PIR, secret sharing and garbling at once

Status. Statement: AI-written from Benny Applebaum and Oded Nir, *Advisor-Verifier-Prover Games and the Hardness of Information Theoretic Cryptography*, IACR ePrint 2023/1378. Not yet checked by a human. The source poses this as Hypothesis 1.2 and offers it “both as a working hypothesis or as an (ambitious) target for future works.” It proves (Theorem 1.3) that the hypothesis implies super-polynomial lower bounds on sd-PIR, general secret sharing and fully-decomposable randomized encodings — for each of which no super-linear lower bound is currently known.

Category. *Information-Theoretic Cryptography.*

Conjecture (informal). *Three basic information-theoretic primitives — private information retrieval, general secret sharing, and garbling — have resisted super-linear communication lower bounds for decades, and nothing relates the three questions to each other. The source finds a single proof model that all three reduce to: an advisor sees a function and hands the verifier a secret advice string; the verifier then gets an input, sends one query to an untrusted prover, and must decide. The hypothesis is that when the prover’s answer is only a constant number of bits, the advice and the query together cannot both be short. Prove that and all three lower bounds follow at once.*

1 The setting

“A major open problem in information-theoretic cryptography is to obtain a super-polynomial lower bound for the communication complexity of basic cryptographic tasks. This question is wide open even for very powerful non-interactive primitives such as private information retrieval (or locally-decodable codes), general secret sharing schemes, conditional disclosure of secrets, and fully-decomposable randomized encoding (or garbling schemes). In fact, for all these primitives we do not even have superlinear lower bounds.” — the source’s abstract

The state of the art is worth stating precisely, because it is the measure of how ambitious the hypothesis is. For sd-PIR — constant number of servers, each answering a constant number of bits, minimizing the client’s message — the best upper bounds are $2^{\tilde{O}(\sqrt{n})}$ and a lower bound of Cn for a constant $C > 1$ has been known since the late 1990s; even improving the constant in the basic case of three queries and binary responses requires highly non-trivial techniques. For secret sharing over n -bit predicates, the best upper bound is $(3/2)^{n+o(n)}$ and the best lower bound, from the mid 1990s, is $\Omega(n/\log n)$. So in each case the gap between what is proved and what is believed is exponential, and the two questions are not known to relate to each other: a lower bound for secret sharing does not imply one for PIR, or vice versa.

The source’s contribution is a single model that all of them reduce to.

2 The model

Definition 1 (AVP game, [AN23, §1.1]). An *Advisor-Verifier-Prover* protocol has an offline and an online phase.

Offline. An arbitrary function $f : \{0,1\}^n \rightarrow \{0,1\}$ is selected and delivered only to the advisor, who computes a randomized advice a . The advice is delivered privately to the verifier, and the advisor leaves.

Online. The verifier receives an input $x \in \{0,1\}^n$ and sends a single query $b = b(a, x)$ to an untrusted prover. The prover, who holds x and f , tries to convince the verifier that $f(x) = 1$ by sending a single message $c = c(b, x, f)$. The verifier accepts or rejects by computing a predicate over a, x, c and its private random tape.

The advice is hidden from the prover, and no re-usability is required: soundness is measured with respect to a single use of the system over a fresh advice. All parties are computationally unbounded. The *communication complexity* is $|a| + |b| + |c|$, and the target is perfect completeness with soundness error $1/2$.

Definition 2 (Laconic prover). A prover is *laconic* if it sends only C bits, where C is a constant that does not grow with n .

3 The conjecture

Conjecture 1 (AVPs are Lengthy, [AN23, Hypothesis 1.2]). *Every prover-laconic AVP protocol that works for the class of n -bit predicates must have total communication complexity that is super-polynomial in n .*

Theorem 1 (What it buys, [AN23, Theorem 1.3]). *Under Conjecture 1, sd-PIR over a 2^n -bit database, general secret sharing over n -bit predicates, and fully-decomposable randomized encodings over n -bit predicates all have communication cost that grows super-polynomially in n .*

Remark 1 (The hypothesis is not vacuous, and the gap it must cross). The transforms proving Theorem 1 run the other way too, and as a by-product yield a non-trivial AVP with sub-exponential communication complexity $2^{\tilde{O}(\sqrt{n})}$ and a laconic prover who communicates a single bit. So prover-laconic AVPs at sub-exponential cost exist, and Conjecture 1 is the claim that they cannot be pushed down to polynomial. The quantity to be ruled out sits between $\text{poly}(n)$ and $2^{\tilde{O}(\sqrt{n})}$, which is exactly the shape of the gap in each of the three primitives.

Remark 2 (Why laconicity is the crux). Without the restriction on the prover the model is easy: the source’s own Example 1.1 exhibits a non-laconic AVP with polynomial communication, rooted in the PIR literature and derivable from an n -server PIR protocol. So the hypothesis is false without “prover-laconic”, and any attack on it must use the constant bound on $|c|$.

Remark 3 (What the source proves toward it). As a first step it shows that simple counting-based lower bounds adapt to this model, re-deriving in a unified way several existing bounds that had been proved separately for each primitive. It records the limitation plainly: “Unfortunately, we do not get the best-known lower bounds for any of the above primitives.” So the model is confirmed to support the classical arguments, and confirmed not to improve on them yet.

Remark 4 (A hypothesis, or a target). The source is explicit that Conjecture 1 “can be used both as a working hypothesis or as an (ambitious) target for future works.” It predicts nothing about which way it goes, and neither does this statement. A refutation — a prover-laconic AVP for all n -bit predicates at polynomial communication — would be at least as interesting as a proof, since it would give a genuinely new protocol in a model that generalizes three well-studied primitives at once.

Remark 5 (Related models it is not). The source compares AVPs with one-way communication complexity and with Yao’s model, and notes one difference that could otherwise confuse a reader: there the goal is to compute f , whereas here the goal is to certify that $f(x) = 1$. The difference does not affect asymptotic complexity, since one can reduce computation to certification by running two copies of an AVP protocol concurrently, once for f and once for $\neg f$.

4 Bibliography

- [AN23] Benny Applebaum and Oded Nir. *Advisor-verifier-prover games and the hardness of information theoretic cryptography*. IACR ePrint 2023/1378. The source. The AVP game is described in Section 1.1 on page 3, Hypothesis 1.2 and Theorem 1.3 are on page 4, and the discussion of what is proved toward the hypothesis and the $2^{\tilde{O}(\sqrt{n})}$ AVP is on page 5.
- [Ito87] Mitsuru Ito, Akira Saito and Takao Nishizeki. *Secret sharing scheme realizing general access structure*. Electronics and Communications in Japan, 1989 (earlier version in Globecom 1987). The origin of the secret-sharing question whose worst-case share size the hypothesis would lower-bound.
- [AN21] Benny Applebaum and Oded Nir. *Upslices, downslices, and secret-sharing with complexity of 1.5^n* . CRYPTO 2021. The $(3/2)^{n+o(n)}$ upper bound the source quotes as the state of the art for general secret sharing.