CONJURA · OPEN PROBLEM

The Price in Coins of Using Few Random Sources

Cutting an n-party protocol down to t coin-tossing parties should cost a factor $\Theta(t)$ in total randomness — and no more

Status. Statement: AI-written from Geoffroy Couteau and Adi Rosén, *Random Sources in Private Computation*, IACR ePrint 2023/074, ASIACRYPT 2022. Not yet checked by a human. The source states this as its Conjecture 13 and calls proving it an interesting open question. It is unproved in both directions: neither the $O(t \cdot R_t)$ upper bound nor the $\Omega(t \cdot R_t)$ lower bound is known in general.

Category. *Information-theoretic secure computation; randomness complexity.*

Conjecture (informal). *Information-theoretically private multi-party computation needs random coins, and coins are expensive: a party that can toss them needs a well-calibrated physical source, so it is worth asking how few of the n parties need one. The source settles that count exactly — for any deterministic functionality and any $t < n/2$, exactly t parties need to toss coins, even though the adversary may corrupt all t of them. But its protocols pay for the small count: they use about t times as many coins in total as protocols in which everybody may toss. The conjecture is that this exchange rate is the true one. For functionalities complex enough to need an honest majority, driving the number of coin-tossing parties down to the minimum costs a factor $\Theta(t)$ in total randomness — necessary as well as sufficient.*

1 The setting

Fix n parties $P_1, \ldots, P_n$, each holding a private input, computing an n-party functionality $\mathcal{F}$ over perfectly private authenticated channels. A protocol is *t-private* if it is perfectly correct and the joint view of any coalition of at most t parties reveals nothing beyond that coalition's own inputs and outputs; adversaries are semi-honest and statically chosen.

Call a party a *source* if it tosses coins during the protocol; the

remaining parties are deterministic and act only on what they receive. Two resources are then being counted at once, and the conjecture is about the exchange rate between them: the number of sources, and the *randomness complexity*, the total number of coins tossed by all parties in the worst case.

Theorem 1 (The exact source count, [CR23, Thm. 1]). *For every deterministic n-party functionality $\mathcal{F}$ and every $t < n/2$, there is a t-private protocol computing $\mathcal{F}$ in which only some fixed set of t parties tosses coins. For randomized functionalities $t+1$ sources are necessary and sufficient.*

Theorem 1 matches a lower bound of Kushilevitz and Mansour, so the count is settled. What it costs is not. The source's protocols combine an outer GMW protocol with an inner BGW protocol, isolating the t sources in sub-computations that touch no input; each coin a deterministic party would have tossed is instead assembled from coins sent by *all* t sources, because the deterministic party cannot know which of them the adversary has corrupted. That is where the factor t enters.

2 Notation and parameters

- n is the number of parties, t the corruption threshold.
- A *source* is a party that tosses coins. The *t-source* randomness complexity of $\mathcal{F}$, written R_t^{src}, is the least randomness complexity of any t-private protocol for $\mathcal{F}$ that uses at most t sources.
- R_t is the least randomness complexity of any t-private protocol for $\mathcal{F}$, with no restriction on how many parties toss coins.
- A functionality is called *complex* here exactly when it cannot be t-privately computed for $t \geq n/2$ — the same dividing line as the classification of Chor and Kushilevitz, and the line at which BGW needs an honest majority.

3 The conjecture

Conjecture 1 ([CR23, Conjecture 13]). *Let $\mathcal{F}$ be an n-party functionality that cannot be t-privately computed for $t \geq n/2$. For any $t < n/2$, let R_t be the randomness complexity of t-privately computing $\mathcal{F}$ (with any number of sources). Then the t-source randomness complexity of $\mathcal{F}$ is $\Theta(t \cdot R_t)$.*

Remark 1 (How the source states it). Conjecture 1 is quoted verbatim from the source, page 16. Its own framing, page 4: "Our conjecture states that, for such functionalities, a $\Theta(t)$ blowup in randomness complexity is necessary and sufficient to minimize the number of random sources. We view this conjecture as an interesting open question." And page 15: "We view the proof of Conjecture 13 an interesting open question."

Remark 2 (Both halves are open, and they are different problems). The upper bound $R_t^{\text{src}} = O(t \cdot R_t)$ says a factor t always suffices. The source's constructions realise it in the cases they treat, but not as a generic transformation from an arbitrary randomness-optimal t-private protocol.

The lower bound $R_t^{\text{src}} = \Omega(t \cdot R_t)$ says a factor t is unavoidable, and it is the harder half, for a reason the source is candid about: unconditional randomness lower bounds are scarce. "Characterizing the minimal amount of randomness required for securely computing a functionality is non-trivial in general, and indeed, no such general characterization is known. We expect that relating the randomness complexity to the number of sources might be of comparable difficulty in general."

Remark 3 (The argument behind the guess, and where it is only an intuition). The source gives a reason to expect the factor t, and flags its status: "We warn the reader that what follows is a purely intuitive reasoning: our goal here is to develop an intuition about which conjecture can reasonably be expected."

The reasoning. In any protocol for a nonlinear functionality some parties — call them sensitive — receive messages that hide intermediate values of the computation behind random coins. Suppose the adversary corrupts some sources together with a sensitive party. When a deterministic party sends a sensitive message, the

coins masking it must be uncorrupted, and they can come only from sources. But the deterministic party cannot know which sources are corrupted, so it appears to have no option but to aggregate coins from all t of them — one coin tossed becomes t coins sent. The reasoning visibly breaks for linear functionalities, where shares can be manipulated locally and no such message exists, which is why the conjecture is restricted to functionalities needing an honest majority.

Remark 4 (The one piece of contrary evidence, and why it does not refute the conjecture). XOR looks like a counterexample and is not. The best known upper bound on the randomness complexity of t-private n-party XOR, due to Kushilevitz and Mansour, is achieved using exactly t sources — minimal sources at no randomness cost at all. The source notes this: it "seemingly contradicts the intuition that t-source private computation should require more randomness than private computation without limitations on the number of sources". XOR is linear, so it falls outside the conjecture's hypothesis, and the apparent contradiction is precisely what the honest-majority restriction is there to exclude.

The source also gives a positive result on the other side, for a constant t, where the conjecture predicts nothing: a 1-private n-party AND protocol using a *single* source and 6 random bits, improving the two-source, 8-bit protocol of Kushilevitz, Ostrovsky, Prouff, Rosén, Thillard and Vergnaud in both counts at once. "Note that t being a constant captures a setting where matching the best known randomness complexity would not contradict Conjecture 13."

Remark 5 (What a resolution has to name). Θ hides two claims and a partial answer settles one of them, so a resolution should say which. Also worth stating explicitly is the class it covers: the hypothesis is a property of $\mathcal{F}$ (no t-private protocol for $t \geq n/2$), not of a particular protocol, and a construction beating $\Omega(t \cdot R_t)$ for one such $\mathcal{F}$ refutes the conjecture outright.

4 Bibliography

[CR23] Geoffroy Couteau and Adi Rosén. *Random sources in private computation.* IACR ePrint 2023/074; ASIACRYPT 2022, LNCS, Springer.

The source. The conjecture is stated on page 16 as Conjecture 13 and motivated on pages 4 and 15; the source-count characterisation is Theorem 1 and Table 1; the single-source AND protocol is Theorem 14.

[KM96] Eyal Kushilevitz and Yishay Mansour. *Randomness in private computations.* 15th ACM PODC, pages 181–190, ACM, August 1996. Cited by the source for the lower bound of t sources that Theorem 1 matches, and for the best known randomness upper bound for t-private XOR — which uses exactly t sources.

[CK91] Benny Chor and Eyal Kushilevitz. *A zero-one law for Boolean privacy.* SIAM Journal on Discrete Mathematics, 4(1):36–47, 1991. The dichotomy the conjecture's hypothesis tracks: which Boolean functions can be n-privately computed and which only t-privately for $t < n/2$.

[KOP+19] Eyal Kushilevitz, Rafail Ostrovsky, Emmanuel Prouff, Adi Rosén, Adrian Thillard and Damien Vergnaud. *Lower and upper bounds on the randomness complexity of private computations of AND.* TCC 2019, Part II, LNCS 11892, pages 386–406, Springer, December 2019. The state of the art the source improves on: 8 random bits from two sources for 1-private n-party AND, with a lower bound of more than one bit.

[BGW88] Michael Ben-Or, Shafi Goldwasser and Avi Wigderson. *Completeness theorems for non-cryptographic fault-tolerant distributed computation.* 20th ACM STOC, pages 1–10, ACM Press, May 1988. The protocol whose AND-gate sub-protocol is both the honest-majority requirement the conjecture's hypothesis refers to and the component that carries the t-fold randomness blowup.

[GR03] Anna Gál and Adi Rosén. *Lower bounds on the amount of randomness in private computation.* 35th ACM STOC, pages 659–666, ACM Press, June 2003. Cited by the source among the randomness lower bounds that exist; the scarcity of such bounds is why the conjecture's lower-bound half is the hard one.