

Statistically Secure Robust Additive Randomized Encodings

The half of the source's main question that survives its refutation

Status. Statement: AI-written from Shai Halevi, Yuval Ishai, Eyal Kushilevitz and Tal Rabin, *Additive Randomized Encodings and Their Applications*, IACR ePrint 2023/870. Not yet checked by a human. The source asks whether all functions admit a statistically secure ARE, robust or non-robust, and strongly conjectures the answer is negative. The *non-robust* half was subsequently refuted: Bitansky, Erabelli, Garg and Ishai construct statistical AREs for all finite functions. The *robust* half is open, and is re-posed as an open question by that same later work.

Category. *Information-Theoretic Cryptography.*

Conjecture (informal). *In an additive randomized encoding each party turns its input into a single element of an abelian group, and the sum of all the elements is supposed to reveal the function's output and nothing else. That is the non-robust notion, and it assumes the evaluator is on its own. A robust ARE assumes the worst instead: the evaluator may collude with some of the parties, in which case it can subtract their encodings and see the sum of the honest ones alone. Since that sum determines the residual function on the honest inputs, the definition concedes the residual function to the simulator and asks that nothing more leak. The source introduced both notions and conjectured that neither admits a statistically secure scheme for all functions. For the non-robust notion that conjecture has since been refuted. For the robust notion the question stands.*

1 The setting

An ARE for a multiparty f over an abelian group $\mathbb{G}$ is a triple (Setup, Enc, Dec): party i maps its input to a group element $\hat{x}_i = \text{Enc}(pp, i, x_i)$, and Dec recovers $f(x_1, \dots, x_n)$ from the sum. Security asks that the sum be simulatable from the output alone. The source, which introduced the notion, also introduced its robust

strengthening, and the reason is structural: a collusion between the evaluator and the parties outside a set H can subtract their own encodings from $\widehat{y}$ and obtain $\widehat{y}_H = \sum_{i \in H} \widehat{x}_i$, the sum over the honest parties. That quantity necessarily determines the residual function $f_{H,x}$ obtained by fixing the honest inputs, so a robust definition cannot hide it; instead it hands the simulator oracle access to $f_{H,x}$ and demands that nothing beyond it leak (Definition 2).

The source’s own results split by accuracy. For *perfect* security it gives a full two-party characterization: AREs exist only for degenerate functions. For *statistical* security it constructs AREs for capped sum (with payloads) and, via that, for multiplication modulo an arbitrary integer, and proves there is no perfectly secure ARE for capped sum — a provable separation between perfect and statistical ARE. It notes that all of its information-theoretic constructions are in fact robust. For *computational* security all polynomial-time functions have an ARE under a bilinear-group assumption, and robust computational ARE for general functions over large domains implies obfuscation, with a converse via resettable MPC.

Its Section 1.2 then records: “The main open question is whether all functions admit a statistically secure (robust or non-robust) ARE. This is open even for very simple functions, such as equality of two inputs from the domain $\{0, 1, 2\}$, which can be shown to be complete for non-robust ARE (Theorem 5.10). We strongly conjecture that the answer is negative.” The closest it gets to a proof is a negative result under an alternative definition using ℓ_2 rather than ℓ_1 distance, documented in its Appendix A.

That conjecture is now half wrong. Bitansky, Erabelli, Garg and Ishai [BEGI25] construct, for every multiparty f over a finite domain and every $\varepsilon > 0$, an ARE with statistical correctness and security errors at most ε — so statistical *non-robust* AREs exist for all finite functions, and the equality-over- $\{0, 1, 2\}$ case the source singled out is settled the other way. Their own open-questions list then re-poses the robust case, noting that a positive answer even for finite 3-party functions would settle the main open question about multi-party randomized encodings.

2 Notation and parameters

- f is an n -party function; the statement below concerns *finite* f , i.e. constant-size domains, which is the regime where the question is information-theoretic at all.
- $\mathbb{G}$ is a finite abelian group, described by the public parameters pp ; λ is the security parameter, ℓ the input length.
- $H \subseteq [n]$ is the set of honest parties; $x = (x_i : i \in H)$ their inputs.
- $\Pi_H(1^\lambda, x_1, \dots, x_n)$ denotes $(pp, \hat{y}_H)$ where $\hat{y}_H = \sum_{i \in H} \hat{x}_i$.
- $f_{H,x}$ is the residual function on the $n - |H|$ remaining inputs, obtained by fixing x_i for $i \in H$.
- δ is the statistical security error; SD is statistical distance.

3 The conjecture

Definition 1 (ARE, [HIKR23, Definitions 3.1 and 3.3]). An ARE for f is $\Pi = (\text{Setup}, \text{Enc}, \text{Dec})$ where $\text{Setup}(1^\lambda, 1^n, 1^\ell) \rightarrow pp$ fixes an abelian group $\mathbb{G}$, $\text{Enc}(pp, i, x_i) \rightarrow \hat{x}_i \in \mathbb{G}$, and $\text{Dec}(pp, \hat{y}) \rightarrow y$. It is correct with error ε if $\text{Dec}(pp, \sum_i \hat{x}_i) = f(x_1, \dots, x_n)$ with probability at least $1 - \varepsilon$, for all inputs. It is *statistically secure* if for some negligible δ there is a simulator Sim with $\text{SD}(\text{Sim}(1^\lambda, 1^n, 1^\ell, f(x_1, \dots, x_n)), \Pi(1^\lambda, x_1, \dots, x_n)) \leq \delta(\lambda)$ for all inputs.

Definition 2 (Robust ARE, [HIKR23, Definition 3.4]). Π is *statistically robust* if for some negligible δ there is a simulator Sim with access to a residual-function oracle such that for all λ, n, ℓ , all $H \subseteq [n]$ and all inputs $x = (x_i : i \in H)$,

$$\text{SD}\left(\text{Sim}^{f_{H,x}}(1^\lambda, 1^n, H, 1^\ell), \Pi_H(1^\lambda, x)\right) \leq \delta(\lambda).$$

Conjecture 1 (Statistically secure robust ARE for finite functions). *Every finite multiparty function f admits an ARE (Definition 1) that is statistically correct and statistically robust (Definition 2), with both errors negligible.*

Remark 1 (Which half of the source’s question this is, and why the other half is gone). The source’s main open question covers robust and non-robust statistical ARE together, and conjectures both are impossible. The non-robust half is now settled affirmatively by [BEGI25], against that conjecture, for every function over a finite domain. Conjecture 1 is therefore the surviving half, stated at the robust notion, and it is stated as an achievability claim rather than as the source’s negative prediction — because the one prediction of that pair which has been tested turned out to be wrong. A refutation is equally a resolution, and the source’s own ℓ_2 -distance negative result in its Appendix A is the closest thing to a route toward one.

Remark 2 (Why finiteness is not a technicality). For general functions over large input domains, robust ARE implies obfuscation — the source proves this, and the reason is that the simulator’s access to the residual function is exactly an obfuscation of it. So there is no information-theoretic question to ask outside the finite regime, and Conjecture 1 is restricted to finite f for that reason and not for convenience. It matches how [BEGI25] poses the robust question, “in the case of finite functions”.

Remark 3 (What a positive answer would settle elsewhere). By [BEGI25], a positive answer even for finite 3-party functions would settle the main open question about multi-party randomized encodings [ABT21], and a positive answer for all finite functions would settle the main open question about best-possible information-theoretic MPC [HIKR18]. The source itself makes the second connection when it says that information-theoretic robust ARE implies best-possible information-theoretic MPC.

Remark 4 (Neighbouring questions in the source’s list). The source also asks: extending its perfect-security characterization beyond two parties; ruling out an ARE for OR with perfect correctness; which assumptions computational ARE really needs (whether other public-key assumptions such as DDH or LWE suffice, and whether public-key cryptography is needed at all — since answered in part by [BEG25], which builds computational ARE from any public-key encryption); constructing robust ARE for constant-size functions from iO and standard assumptions rather than ideal obfuscation; and applying robust ARE to differential privacy in the robust shuffle model. None of those is this statement.

4 Bibliography

- [HIKR23] Shai Halevi, Yuval Ishai, Eyal Kushilevitz and Tal Rabin. *Additive randomized encodings and their applications*. IACR ePrint 2023/870. The source. Definitions 3.1, 3.3 and 3.4 are on pages 10–11; the summary of what is left open is on page 5 and the open-questions list in Section 1.2, page 7.
- [BEGI25] Nir Bitansky, Saroja Erabelli, Rachit Garg and Yuval Ishai. *Shuffling is universal: statistical additive randomized encodings for all functions*. IACR ePrint 2025/1442. Refutes the non-robust half of the source’s conjecture, and re-poses the robust case for finite functions. [UNVERIFIED] — it postdates the source, so it is not in the source’s reference list; it was checked against the paper itself in the AICR harvest that produced c/0073.
- [BEG25] Nir Bitansky, Saroja Erabelli and Rachit Garg. *Additive randomized encodings from public key encryption*. IACR ePrint 2025/104. Builds computational ARE from any public-key encryption, bearing on the source’s question about which assumptions are needed. [UNVERIFIED] — postdates the source; read in the reference list of [BEGI25], not of the source.
- [ABT21] Benny Applebaum, Zvika Brakerski and Rotem Tsabary. *Perfect secure computation in two rounds*. SIAM Journal on Computing, 50(1):68–97, 2021. The multi-party randomized encoding question a positive answer for finite 3-party functions would settle.
- [HIKR18] Shai Halevi, Yuval Ishai, Eyal Kushilevitz and Tal Rabin. *Best possible information-theoretic MPC*. Theory of Cryptography Conference (TCC) 2018, Part II, pages 255–281. The notion the source states information-theoretic robust ARE implies.
- [IK00] Yuval Ishai and Eyal Kushilevitz. *Randomizing polynomials: a new representation with applications to round-efficient secure computation*. 41st Annual Symposium on Foundations of Computer Science (FOCS), 2000. Introduces randomized encodings, of which ARE is an instance, and poses the degree-2 statistical RE question the source flags as a separate 20-year-old open problem.