CONJURA · OPEN PROBLEM

Round-by-Round Soundness Amplification Under Parallel Repetition

r rounds, error ϵ: does $(k \cdot r)$-fold repetition give ϵ^k?

Status. Statement: AI-written from Dor Bitan, Zachary DeStefano, Shafi Goldwasser, Yuval Ishai, Yael Tauman Kalai and Justin Thaler, *Sum-Check Protocol for Approximate Computations*, IACR ePrint 2025/2152. Not yet checked by a human. Known in the special case of r-round protocols whose standard and round-by-round soundness errors are maximally separated, $\epsilon_{\mathrm{RBR}} \approx \epsilon^{1/r}$; open in general.

Category. *Proof Systems.*

Conjecture (informal). *Round-by-round soundness is the notion that governs how well an interactive proof survives the Fiat–Shamir transform: a doomed transcript must stay doomed, round after round, except with probability ϵ per round. Ordinary soundness amplifies under parallel repetition in the familiar way. Round-by-round soundness is more delicate, because what has to be amplified is a per-round invariant rather than an overall acceptance probability, and the doomed set of the repeated protocol has to be assembled from those of the copies. The conjecture is that $k \cdot r$ parallel copies of an r-round protocol with round-by-round error ϵ have round-by-round error at most ϵ^k. If true, it means Fiat–Shamir security can be bought by repetition alone — from $\log(1/\epsilon)$ bits to $k \log(1/\epsilon)$ bits — without raising the numerical precision of any single instance, which is what the source needs in its approximate setting.*

1 The setting

The source generalizes the sum-check protocol to approximate claims $\sum_{x \in \{0,1\}^v} g(x) \approx H$ over an integral domain, with a tunable error parameter δ; its soundness degrades gracefully in δ/Δ, where Δ is the error in the prover's initial claim. Because the intended use is a Fiat–Shamir'ed non-interactive argument, the quantity that matters is not standard soundness but round-by-round soundness:

by the standard transfer [BSCS16, CCH+19, BGTZ25], an r-round public-coin proof with round-by-round error ϵ_{RBR} Fiat–Shamirs to soundness error $O(Q\epsilon_{\mathrm{RBR}} + Q^2 2^{-\kappa})$ against a prover making Q random-oracle queries.

The source's Theorem 4.8 gives, for its protocol instantiated over $\mathbb{C}$ with the absolute-value metric where S is a set of roots of unity, round-by-round soundness error at most $d/|S| + 2\sqrt[d]{\delta/\Delta}$. Its own summary notes that the resulting regime is roughly $\epsilon^{2/r}$, intermediate between standard soundness and the tightest round-by-round behaviour — a phenomenon it argues looks inherent to approximate computation, since a cheating prover can make an inaccurate answer look slightly less erroneous each round.

Immediately after that theorem the source records a conjecture about round-by-round soundness in general, not specific to its approximate protocol: "We conjecture the following general amplification property: Let Π be an r-round interactive protocol with round-by-round (RBR) soundness error ϵ. Then, $(k \cdot r)$-fold parallel repetition of Π results in a protocol with RBR soundness error at most ϵ^k." It states the consequence it cares about — that $(k \cdot r)$-fold parallel repetition boosts the Fiat–Shamir security level from $\log(1/\epsilon)$ bits to $k \log(1/\epsilon)$ bits, without any increase in the precision of individual instances — and it names the one case where the conjecture is known: r-round protocols whose standard soundness error ϵ and round-by-round soundness error ϵ_{RBR} are maximally separated, i.e. $\epsilon_{\mathrm{RBR}} \approx \epsilon^{1/r}$, by [CCH+18, Corollary 5.7].

That special case is where the conjecture's content is thinnest: maximal separation is precisely the regime in which round-by-round soundness is as weak as it can be relative to standard soundness, so the amplification has the most room. What is open is every other regime, including the $\epsilon^{2/r}$ intermediate one the source's own protocol lands in.

2 Notation and parameters

- $\Pi = (P, V)$ is a public-coin interactive proof for a language L, with r rounds.
- ϵ is the round-by-round soundness error of Π; ϵ_{RBR} is used where the contrast with standard soundness error matters.
- $k \geq 1$ is the repetition multiplier; the repeated protocol runs

$k \cdot r$ independent copies in parallel, the verifier accepting only if all copies accept.

- τ denotes a partial transcript, α a prover message, β a verifier message, and $\tau|\alpha|\beta$ their concatenation.
- Q is a bound on a malicious prover's random-oracle queries and κ the oracle's output length, in the Fiat–Shamir statement.

3 The conjecture

Definition 1 (Round-by-round soundness, [BDGIKT25, Definition 3], after [CCH+19, Hol19]). A public-coin interactive proof $\Pi = (P, V)$ for a language L has *round-by-round soundness error* ϵ_{RBR} if there is a "doomed set" D of partial transcripts such that

(1) if $x \notin L$ then $(x, \emptyset) \in D$;
(2) for every partial transcript τ, prover message α and subsequent verifier message β, if $(x, \tau) \in D$ then $\Pr_\beta[(x, \tau|\alpha|\beta) \notin D] \le \epsilon_{\mathrm{RBR}}$; and
(3) for every complete transcript τ, if $(x, \tau) \in D$ then $V(x, \tau) = \text{reject}$.

Conjecture 1 (RBR amplification under parallel repetition, [BDGIKT25, Remark 3]). *Let Π be an r-round public-coin interactive proof with round-by-round soundness error ϵ (Definition 1). Then for every $k \ge 1$, the $(k \cdot r)$-fold parallel repetition of Π has round-by-round soundness error at most ϵ^k.*

Remark 1 (The consequence the source draws). The source states that Conjecture 1 implies $(k \cdot r)$-fold parallel repetition boosts the Fiat–Shamir security level from $\log(1/\epsilon)$ bits to $k \log(1/\epsilon)$ bits, "without requiring an increase in the precision of individual instances". That last clause is why the conjecture matters in the source's own setting: raising precision is what its approximate protocol would otherwise have to pay to reach a target security level, and repetition would be a way to avoid it.

Remark 2 (What is already known, and why it is the easy case). Conjecture 1 is known when ϵ and the protocol's standard

soundness error are maximally separated, i.e. when $\epsilon_{\mathrm{RBR}} \approx \epsilon_{\mathrm{sound}}^{1/r}$ ([CCH+18, Corollary 5.7], as cited by the source). A resolution has to cover the regimes in between — in particular the $\epsilon^{2/r}$ regime the source's own approximate sum-check protocol occupies, and the regime where the two errors essentially coincide, which is where the classical sum-check protocol sits ($\epsilon_{\mathrm{RBR}} = d/q$ against standard soundness vd/q).

Remark 3 (Both directions, and where a counterexample would live). The statement is a conjecture, so a proof settles it; so does a counterexample, an r-round protocol with round-by-round error ϵ whose $(k \cdot r)$-fold parallel repetition has round-by-round error above ϵ^k for some k. Note that the difficulty is not that soundness fails to amplify — for public-coin proofs ordinary soundness does — but that Definition 1 requires exhibiting a doomed set for the repeated protocol, and the natural candidate (the transcript is doomed if enough copies are doomed) does not obviously satisfy clause (2) with the claimed error. A counterexample would most plausibly be a protocol where the copies' doomed sets can be escaped in a correlated way across rounds.

Remark 4 (The repetition count is $k \cdot r$, not k). The conjecture ties the number of copies to the round count: $k \cdot r$ copies buy ϵ^k, so the exponent gained is the number of copies divided by the number of rounds. A statement with k copies giving ϵ^k would be strictly stronger and is not what the source conjectures; a statement with $k \cdot r$ copies giving $\epsilon^{k/r}$ would be weaker. The factor r is what makes the conjectured bound consistent with the known maximally-separated case, where per-round error ϵ corresponds to overall error about ϵ^r.

4 Bibliography

[BDGIKT25] Dor Bitan, Zachary DeStefano, Shafi Goldwasser, Yuval Ishai, Yael Tauman Kalai and Justin Thaler. *Sum-check protocol for approximate computations.* IACR ePrint 2025/2152. The source. Definition 3 is on page 7; the conjecture is Remark 3, directly after Theorem 4.8, page 22 (PDF page 24).

[CCH+19] Ran Canetti, Yilei Chen, Justin Holmgren, Alex Lombardi, Guy N. Rothblum, Ron D. Rothblum and Daniel Wichs. *Fiat–Shamir: from practice to theory.* 51st Annual ACM SIGACT Symposium on Theory of

Computing (STOC), 2019, pages 1082–1090. The origin of round-by-round soundness, and of the d/q round-by-round bound for classical sum-check.

[Hol19] Justin Holmgren. *On round-by-round soundness and state restoration attacks.* IACR ePrint 2019/1261. The other source of Definition 1.

[CCH+18] Ran Canetti, Yilei Chen, Justin Holmgren, Alex Lombardi, Guy N. Rothblum and Ron D. Rothblum. *Fiat–Shamir from simpler assumptions.* IACR ePrint 2018/1004. Its Corollary 5.7 gives the one case in which Conjecture 1 is known: r-round protocols with maximally separated soundness and round-by-round soundness errors.

[BSCS16] Eli Ben-Sasson, Alessandro Chiesa and Nicholas Spooner. *Interactive oracle proofs.* 14th International Conference on Theory of Cryptography (TCC), Part II, LNCS 9986, 2016, pages 31–60. One of the works the Fiat–Shamir transfer theorem follows from.

[BGTZ25] Alexander R. Block, Albert Garreta, Pratyush Ranjan Tiwari and Michał Zając. *On soundness notions for interactive oracle proofs.* Journal of Cryptology, 38(1):4, 2025. The third of the works the source's Fiat–Shamir transfer theorem (its Theorem 2.2) follows from.

[LFKN92] Carsten Lund, Lance Fortnow, Howard Karloff and Noam Nisan. *Algebraic methods for interactive proof systems.* Journal of the ACM, 39(4):859–868, 1992. The classical sum-check protocol the source generalizes.