AICR · OPEN PROBLEM

A Dream Direct Product Theorem for a Moderately Hard Function

exponential multi-instance hardness from a standard fine-grained hypothesis

Status. Statement: AI-written from Damiano Abram, Marshall Ball, Juan Garay and Aggelos Kiayias, *Permissionless Consensus from a Common Random String*, IACR Cryptology ePrint Archive, 2026 (abridged version to appear in Proc. CRYPTO 2026), not yet checked by a human. No case is settled: the source reduces its consensus protocol to a moderately hard function with these parameters and then leaves the construction of such a function open, offering no candidate.

Category. *Fine-Grained Hardness & Proofs of Work*

Conjecture (informal). *A proof of work is a puzzle that costs a lot of computation to solve and almost nothing to check. Building permissionless consensus without a random oracle needs a puzzle with an unusually strong amplification property: an adversary handed many independent puzzles, and given only about as much time as it takes to solve them one after another honestly, should succeed on all of them with probability at most its single-puzzle success probability raised to the number of puzzles, with no slack in the exponent. This matters because the machinery that lets parties agree on puzzles without a beacon loses a factor exponential in the number of participants, so only an exponentially small multi-instance success probability survives the loss. The obvious workaround — enlarge the security parameter until the puzzle is hard enough — is unavailable, because it also makes each honest party's work grow with the number of participants. Known puzzles built from fine-grained complexity assumptions do not reach the required level, while the ad hoc puzzles used in practice probably do. Since a strong enough succinct argument system turns any hard-to-compute function into an efficiently checkable puzzle, the whole gap reduces to a question in hardness amplification: exhibit a function, together with a standard fine-grained complexity hypothesis, for which computing many independent instances quickly is exactly as hard as the product of the individual instances.*

1 The setting

Permissionless consensus asks a set of parties with no PKI, no authenticated channels and no knowledge of each other to reach Byzantine agreement, tolerating an adversary that controls a fraction of the collective computational power. Every deployed protocol in this space, starting with Nakamoto's, relies on the random oracle heuristic or on ad hoc assumptions. Ball, Garay, Hall, Kiayias and Panagiotakos [1] gave the first provably secure protocol outside the random oracle model, but in the *random beacon model*: all parties are handed fresh public randomness at the start of each voting phase, which is what lets them agree on which proof-of-work challenges are legitimate.

The present paper removes the beacon. Its route is to let the parties generate the challenges themselves, in one round, with a

distributed sampler [3], and then to verify a vote only against a challenge derived from the verifier's own sampler message. Because a single honest sampler message must be reused across many differently labelled executions, the paper introduces d-wise independent hardness-preserving distributed samplers and builds them from DDH and from LWE. Hardness preservation is a game-based, non-tight guarantee: an attack that succeeds with probability p_{real} against the sampler's output translates into an attack succeeding with probability at least $p_{\mathsf{real}}/\eta - \nu$ against a genuinely random sample. The paper shows the loss η must be exponential in the number d of reused executions, and its constructions achieve $\eta(t,\lambda) = 2^{t \cdot n^{\gamma'}}$ for a constant $\gamma' < 1$, where n is the output length.

Consequently the underlying proof of work must be secure against a multiplicative loss of that size. Corollary 3 of the paper composes the sampler with a $(d_0, d_1, \beta, \epsilon)$-moderately hard proof of work whose multi-instance soundness is $\epsilon(t,\lambda) = 2^{-t \cdot n^{\gamma}}$ for a constant $\gamma > 0$, and obtains a multi-verifier signature of work, from which the paper's consensus protocol follows. Two things block the instantiation. First, existing proofs of work built from fine-grained hardness assumptions [1] do not reach subexponential soundness. Second, complexity leveraging is ruled out: raising the security parameter of the puzzle in proportion to d also raises the honest solving cost in proportion to d, so every honest vote would cost work proportional to the total number of participants.

What the paper does show is that efficient verifiability is not the obstacle: any function that is moderately hard with the parameters of Corollary 3, even if it is not efficiently verifiable, becomes a proof of work when compiled with a subexponentially sound SNARG for P, which follows from subexponentially secure LWE by the work of Choudhuri, Jain and Jin [2]. So the entire remaining gap is a hardness amplification statement about a plain function, and the paper closes by asking for it: a tight "dream" strong direct product theorem for some function, assuming a standard fine-grained hypothesis such as NSETH — which the paper invokes by name on p. 23 without defining or citing it. Settling it would put permissionless consensus in the plain CRS model on standard complexity conjectures rather than on an assumed puzzle.

2 Notation and parameters

- $\lambda \in \mathbb{N}$ is the security parameter and $s \in \mathbb{N}$ is a *time parameter*; both are given to all algorithms in unary. Running time is counted in the Word RAM model with $O(\log \lambda)$-bit words; this is the ambient computational model that the source fixes in its Section 4.1 (p. 16) for protocol executions, and not part of Definition 1 below.
- $n = n(\lambda)$ is the challenge length; challenges are strings in $\{0,1\}^n$.
- Gen is a PPT *challenge sampler*: $\mathsf{Gen}(1^\lambda)$ outputs $z \in \{0,1\}^n$.
- $f = \{f_{s,\lambda}\}_{s,\lambda \in \mathbb{N}}$ is a family of functions $f_{s,\lambda} : \{0,1\}^n \to \{0,1\}^*$, and $r(s,\lambda)$ is the number of steps needed to evaluate $f_{s,\lambda}$ on one input (the honest solving cost).
- $\beta = \beta(s,\lambda) \in [0,1]$ is the *slowdown* factor: an adversary asked to solve t challenges is granted $t \cdot \beta(s,\lambda) \cdot r(s,\lambda)$ steps. Larger β means a tighter hardness guarantee.
- $d_0 = d_0(\lambda)$ and $d_1 = d_1(\lambda)$ are polynomials bounding, respectively, the number of challenges handed to the adversary and the number of oracle calls it may make.
- O_f is the oracle that on input $z \in \{0,1\}^n$ returns $f_{s,\lambda}(z)$.
- $\epsilon = \epsilon(t,\lambda) \in [0,1]$ bounds the adversary's success probability on t challenges, and $\gamma > 0$ is the constant governing it.
- N denotes the number of Boolean variables of a formula; it is used only in the statement of NSETH.
- $[t]$ denotes $\{1,2,\ldots,t\}$.

Definition 1 (Moderately hard function). Let d_0, d_1 be polynomials, let $\beta : \mathbb{N} \times \mathbb{N} \to [0,1]$ and $\epsilon : \mathbb{N} \times \mathbb{N} \to [0,1]$ be functions, and let (Gen, f) be a challenge sampler together with a function family as above, with evaluation cost $r(s,\lambda)$. We say (Gen, f) is $(d_0, d_1, \beta, \epsilon)$-*moderately hard* if for every $t \le d_0(\lambda)$ and every pair of PPT algorithms $\mathcal{A} = (\mathcal{A}_0, \mathcal{A}_1)$ such that $\mathcal{A}_1$

runs in at most $t \cdot \beta(s,\lambda) \cdot r(s,\lambda)$ steps and makes at most $d_1(\lambda)$ queries to O_f, none of which lies in $\{z_1,\ldots,z_t\}$, it holds for all sufficiently large $\lambda, s \in \mathbb{N}$ that

$$\Pr\left[\, \forall h \in [t] : w_h = f_{s,\lambda}(z_h) \,\right] \;\leq\; \epsilon(t,\lambda),$$

where the probability is over the experiment

$$\begin{aligned} \mathsf{aux} &\leftarrow \mathcal{A}_0(1^\lambda, 1^s), \\ z_1,\ldots,z_{d_0(\lambda)} &\leftarrow \mathsf{Gen}(1^\lambda) \quad \text{i.i.d.}, \\ (w_h)_{h\in[t]} &\leftarrow \mathcal{A}_1^{O_f}\left(1^\lambda, 1^s, (z_h)_{h\in[d_0(\lambda)]}, \mathsf{aux}\right). \end{aligned}$$

Remark 1 (Provenance of Definition 1). Definition 1 is this record's adaptation of Definition 14 of the source (p. 27), and departs from it in three ways, all of them the record's own choices rather than the paper's.

- **The oracle restriction.** The clause "none of which lies in $\{z_1,\ldots,z_t\}$" does not appear in Definition 14, which bounds only the number of oracle queries by $d_1(\lambda)$. It is added here because Definition 14 as printed is trivially breakable whenever $d_1 \geq t$: the adversary can simply query the oracle on the challenges it must solve. The analogous restriction does appear in Definition 5 of the source (p. 14), as $m_h \notin Q \cup \{m_j\}_{j\neq h}$. The repair is substantive and should not be attributed to the paper.
- **A plain function rather than a proof of work.** Definition 14 is stated for a proof of work $\mathsf{PoW} = (\mathsf{Setup}, \mathsf{Gen}, \mathsf{Solve}, \mathsf{Verify})$: there is a CRS $\sigma \leftarrow \mathsf{Setup}(1^\lambda, 1^s)$, the oracle O_σ^s returns $\mathsf{Solve}(1^s, \sigma, z)$, and the win condition is $\mathsf{Verify}(\sigma, z_h, w_h) = 1$. The CRS-free, function-only reading used here — oracle O_f, win condition $w_h = f_{s,\lambda}(z_h)$ — follows the source's informal remark on p. 16, that "given any function f that is moderately-hard (Def. 14) with parameters as required by Corollary 3 but not necessarily efficiently verifiable, can be turned into a proof of work via any subexponentially sound SNARG for P", which drops Verify and the CRS.
- **How many challenges are sampled.** Definition 14 sam-

ples only $z_1, \ldots, z_t$ from Gen while handing $\mathcal{A}_1$ the tuple $(z_h)_{h \in [d_0(\lambda)]}$. This appears to be a typo; the record repairs it by drawing all $d_0(\lambda)$ challenges i.i.d. from Gen.

Definition 2 (NSETH). For every $\varepsilon > 0$ there exists $k \in \mathbb{N}$ such that k-TAUT — the problem of deciding whether a given k-DNF formula on N variables is a tautology — cannot be decided by a nondeterministic algorithm running in time $O(2^{(1-\varepsilon)N})$.

Remark 2 (Provenance of Definition 2). Definition 2 is supplied by this record from the standard literature, with [4] as the customary reference; both the wording and the reference are the record's, not the paper's. The source mentions NSETH exactly once, on p. 23, without defining it and without citing anything for it, and its reference list contains no Carmosino et al. entry. The current status of NSETH in the literature — barriers, conditional refutations — has not been checked here.

3 The conjecture

Conjecture 1 (Dream direct product theorem for a moderately hard function). *Assume NSETH (Definition 2). Then there exist a constant $\gamma > 0$, polynomials n, d_0, d_1, a function r, a function $\beta : \mathbb{N} \times \mathbb{N} \to [0,1]$, a PPT challenge sampler* Gen *with output length $n(\lambda)$, and a family $f = \{f_{s,\lambda}\}$ whose evaluation cost is $r(s,\lambda)$, such that* (Gen, f) *is $(d_0, d_1, \beta, \epsilon)$-moderately hard in the sense of Definition 1 with*

$$\epsilon(t,\lambda) \;=\; 2^{-t \cdot n(\lambda)^{\gamma}} \qquad \text{for every } t \le d_0(\lambda).$$

That is, the per-instance bound $2^{-n(\lambda)^{\gamma}}$ amplifies to its exact t-th power over t independent instances, with no loss in the exponent.

Remark 3 (Latitude in the hypothesis). The source writes "assuming, say, NSETH", so NSETH is its example rather than its commitment. Fixing NSETH as in Conjecture 1 narrows the question, and a solver should feel free to substitute another standard fine-grained hypothesis.

Remark 4 (Side conditions the record considered and did not impose). Conjecture 1 constrains neither r nor β beyond their types, and this matches the source: Corollary 3 (p. 16) constrains only $\beta'(s,\lambda,L) \geq \beta(s,\lambda)/2$ for sufficiently large polynomial s, Definition 14 (p. 27) places no relation between the evaluation cost and s, and p. 5 reports $\lim_{\lambda\to\infty} \beta = 0$ for the paper's own protocol. Two further conditions — $s \leq r(s,\lambda) \leq s \cdot \mathrm{poly}(\lambda)$ and $1/\beta(s,\lambda) \leq \mathrm{poly}(\lambda)$ for all s — would make the conjecture harder to satisfy vacuously, but they are nowhere in the source and are recorded here only as the record's own suggestion. A third choice is the record's too: the conjecture states only the upper bound ϵ, and does not additionally demand that the per-instance bound be achievable, which is what "tight" would require.

Remark 5 (What the source establishes). (i) The required soundness must satisfy $\epsilon \leq 2^{-\Omega(d)}$, because the sampler's hardness-preservation loss is $\eta = 2^{\Omega(d)}$. (ii) Complexity leveraging cannot supply it, since raising the puzzle's security parameter in proportion to d raises honest solving time in proportion to d. (iii) Existing proofs of work from fine-grained hardness assumptions [1] fall short of subexponential soundness. (iv) Efficient verifiability is not the bottleneck: a subexponentially sound SNARG for P, available from subexponentially secure LWE via Choudhuri, Jain and Jin [2], upgrades any such function to a proof of work. Nothing later in the paper or its appendices supplies a candidate function; the paper only remarks that the ad hoc puzzles in current use are likely to satisfy the requirement.

Remark 6 (Why this form). Conjecture 1 is about a single function family and a single probability bound. It mentions no distributed sampler, no consensus protocol and no signature scheme: the reduction from the cryptographic setting has already been carried out by the paper, which shows that efficient verifiability can be added for free by a succinct argument. A solver needs only Definition 1 and a fine-grained hypothesis. Conversely, a positive answer would put permissionless consensus in the plain CRS model on a CRS, standard lattice or group assumptions, and a standard fine-grained conjecture — no random oracle, no beacon, no ad hoc puzzle — and dream (lossless) direct product theorems in a fine-grained, nondeterministic-hardness setting are exactly the kind of statement

that resists the standard amplification toolbox.

4 Bibliography

[1] Marshall Ball, Juan A. Garay, Peter Hall, Aggelos Kiayias, and Giorgos Panagiotakos. Towards permissionless consensus in the standard model via fine-grained complexity. CRYPTO 2024, Part II, volume 14921 of LNCS, pages 113–146, Springer, Cham, 2024.

[2] Arka Rai Choudhuri, Abhishek Jain, and Zhengzhong Jin. Non-interactive batch arguments for NP from standard assumptions. CRYPTO 2021, Part IV, volume 12828 of LNCS, pages 394–423, Springer, 2021.

[3] Damiano Abram, Brent Waters, and Mark Zhandry. Security-preserving distributed samplers: How to generate any CRS in one round without random oracles. CRYPTO 2023, Part I, volume 14081 of LNCS, pages 489–514, Springer, Cham, 2023.

[4] Marco L. Carmosino, Jiawei Gao, Russell Impagliazzo, Ivan Mihajlin, Ramamohan Paturi, and Stefan Schneider. Nondeterministic extensions of the strong exponential time hypothesis and consequences for non-reducibility. ITCS 2016 (recalled from memory; not cited in the source paper), 2016. [UNVERIFIED]