CONJURA · OPEN PROBLEM

Improving the Field Size of Large-Field Dot-Product Proofs

Proof length $O(S \log(S/\varepsilon))$, soundness ε: how small may the prime be?

Status. Statement: AI-written from Nir Bitansky, Prahladh Harsha, Yuval Ishai, Ron D. Rothblum and David J. Wu, *Dot-Product Proofs and Their Applications*, IACR ePrint 2024/1138. Not yet checked by a human. The small-field regime is settled: soundness $\Theta(1/\sqrt{q})$ is achieved and proved optimal. In the large-field regime the source achieves soundness ε for every prime $p \geq \widetilde{\Omega}(S^9/\varepsilon^6)$ and states in as many words that it believes the S-dependence is not tight.

Category. *Proof Systems.*

Conjecture (informal). *A dot-product proof is the most laconic proof system imaginable: the statement x and the proof π are vectors over a field, and the verifier is allowed exactly one dot product $\langle q, (x\|\pi)\rangle$ against both of them at once, after which it accepts or rejects. The source constructs such proofs for circuit satisfiability with proof length linear in the circuit size S, avoiding the PCP theorem. To get soundness error ε it needs the field to be large: any prime $p \geq \widetilde{\Omega}(S^9/\varepsilon^6)$ works. The ninth power of the circuit size is what keeps the resulting laconic arguments a proof of concept rather than a usable system — with a standard 256-bit group the source can support statements of around 2^{10} gates at a few bits of soundness. The authors state they believe their analysis is loose and conjecture the dependence on S can be improved; the corresponding improvement is already known for the simpler Hadamard-based linear PCP, by a random-walk argument.*

1 The setting

A dot-product proof (DPP) for a language L over a field $\mathbb{F}$ is a one-query fully linear PCP: the verifier samples a query vector q and an accepting set A, is given oracle access to nothing but the single field element $\langle q, (x\|\pi)\rangle$, and accepts iff that element lies in A. Because the query touches the whole input as well as the proof,

soundness can be required for every $x \notin L$ rather than only for x far from L, unlike a PCP of proximity.

The source gives two constructions. Over constant-size fields (Theorem 1.2) it achieves proof length $S \cdot \mathrm{poly}(q)$ with soundness error $\varepsilon_q = O(1/\sqrt{q})$, and proves a matching $\Omega(1/\sqrt{q})$ lower bound, so that regime is closed. Over large fields (Theorem 1.3) it achieves, for a prime $p \geq \widetilde{\Omega}(S^9/\varepsilon^6)$, a *promise* DPP of proof length $O(S \cdot \log(S/\varepsilon))$ and soundness error ε; proof length can be brought to $O(S)$ at the cost of a worse polynomial dependence of p on S/ε. This large-field construction is what the source's applications to laconic arguments run on, since only there can ε be made negligible.

The field size matters concretely because the argument's proof consists of group elements whose size is $\log p$: at a fixed group size, the S^9 caps the circuit size that can be supported. The source is explicit that the bound is an artefact of its analysis rather than of the construction, on two counts. Of the general fully-linear-PCP-to-DPP compiler it applies, it says: "We believe that the analysis of our compiler is not tight, and leave open the question of improving the bound on p given by Theorem 1.3." And of the concrete efficiency of the resulting argument, in a footnote: "In particular, we conjecture that the dependence on S (i.e., the circuit size) in our current soundness bound can be improved. This was shown in the simpler case of the Hadamard-based LPCP using a random walk argument [BIOW20], and while it seems to heuristically hold also in our setting, the formal analysis is much more challenging."

That is the named obstruction: the random-walk concentration argument that removes the loss for the Hadamard-based linear PCP of Barta et al. [BIOW20] does not obviously transfer to the source's compiler, which is applied to a two-query fully linear PCP implicit in Danezis et al. [DFGK14] rather than to the Hadamard LPCP.

2 Notation and parameters

- $C : \{0,1\}^n \times \{0,1\}^m \to \{0,1\}$ is a Boolean circuit of size S, and $L_C = \{x \in \{0,1\}^n : \exists w \in \{0,1\}^m, C(x,w) = 1\}$.
- p is a prime and $\mathbb{F}_p$ the field; $q = |\mathbb{F}|$ in the small-field discussion.
- $\varepsilon \in (0,1)$ is the soundness error, c the completeness parameter ($c = 1$ is perfect completeness).

- $\widetilde{\Omega}(\cdot)$ hides polylogarithmic factors, as in the source.

3 The conjecture

Definition 1 (Dot-product proof, [BHIRW24, Section 1.1]). A *dot-product proof* for a language $L \subseteq \{0,1\}^n$ over $\mathbb{F}$, with proof length m, is a randomized verifier V outputting a pair (q, A) with $q \in \mathbb{F}^{n+m}$ and $A \subseteq \mathbb{F}$, such that

- (completeness) for every $x \in L$ there is $\pi \in \mathbb{F}^m$ with $\Pr_{(q,A)\leftarrow V}[\langle q, (x\|\pi)\rangle \in A] \geq c$; and
- (soundness) for every $x \notin L$ and every $\pi^* \in \mathbb{F}^m$, $\Pr_{(q,A)\leftarrow V}[\langle q, (x\|\pi^*)\rangle \in A] \leq \varepsilon$.

In a *promise* DPP, soundness is required only for $x \in \{0,1\}^n \setminus L$, with no promise on $\pi^* \in \mathbb{F}^m$. The proof is *efficient* if prover and verifier run in time polynomial in the proof length, including an implicit representation of A when $|\mathbb{F}|$ is super-polynomial.

Conjecture 1 (The field size may be reduced). *There are constants $a < 9$ and $b > 0$ such that the following holds. For every Boolean circuit $C : \{0,1\}^n \times \{0,1\}^m \to \{0,1\}$ of size S, every $\varepsilon > 0$, and every prime $p \geq \widetilde{\Omega}(S^a/\varepsilon^b)$, the language L_C has an efficient promise DPP (Definition 1) over $\mathbb{F}_p$ with perfect completeness, soundness error ε, and proof length $O(S \cdot \log(S/\varepsilon))$.*

Remark 1 (Why the exponent, and not a limit). The source conjectures that "the dependence on S ... can be improved", without naming a target exponent, so Conjecture 1 asks only for some a strictly below the source's 9 while leaving the ε-dependence free. This is the weakest reading under which the source's own conjecture has content, and the honest one: any resolution should be reported with the pair (a, b) it achieves. Two natural stronger targets, neither claimed by the source, are a arbitrarily close to 0 (field size depending on S only polylogarithmically) and the removal of the S-dependence entirely.

Remark 2 (The more ambitious question this is not). The source separately raises "the question of closing the big polynomial gap between the soundness error we achieve and the optimal $O(1/\sqrt{|\mathbb{F}|})$

soundness that we can achieve over small fields", adding that this "might call for an entirely different approach than the one we take here". That is a different statement: Conjecture 1 keeps the source's construction shape and asks for a better analysis of it, whereas closing the gap to $O(1/\sqrt{|\mathbb{F}|})$ in the large-field regime is an achievability question about a possibly different construction. The conjecture stated here is the one the source attaches its own belief to.

Remark 3 (Proof length is held fixed). The $O(S \log(S/\varepsilon))$ proof length in Conjecture 1 is Theorem 1.3's, and it matters that it stays: the source notes the proof length can be reduced to $O(S)$ *at the cost of* a worse polynomial dependence of p on S/ε, so a scheme trading proof length for field size in the other direction would not settle this statement. Likewise the statement is about the promise DPP; the source's non-promise variant costs proof length $O(S + n^2)$ and again a worse dependence of p.

4 Bibliography

[BHIRW24] Nir Bitansky, Prahladh Harsha, Yuval Ishai, Ron D. Rothblum and David J. Wu. *Dot-product proofs and their applications.* IACR ePrint 2024/1138. The source. Theorem 1.2 (small fields) is on page 3, Theorem 1.3 (large fields) and the open question on page 4, and the conjecture on the S-dependence is footnote 6, page 7.

[BIOW20] Ohad Barta, Yuval Ishai, Rafail Ostrovsky and David J. Wu. *On succinct arguments and witness encryption from groups.* CRYPTO 2020, Part I, LNCS 12170, pages 776–806. The source cites this for the random-walk argument that removes the analogous loss for the Hadamard-based linear PCP, and for a laconic argument with a quadratic-size CRS.

[BCIOP22] Nir Bitansky, Alessandro Chiesa, Yuval Ishai, Rafail Ostrovsky and Omer Paneth. *Succinct non-interactive arguments via linear interactive proofs.* Journal of Cryptology, 35(3):15, 2022 (preliminary version in TCC 2013). The linear-only-encryption compiler the source adapts to turn a DPP into a succinct argument.

[DFGK14] George Danezis, Cédric Fournet, Jens Groth and Markulf Kohlweiss. *Square span programs with applications to succinct NIZK arguments.* ASIACRYPT 2014, Part I, LNCS 8873, pages 532–550. The two-query fully linear PCP the source's compiler is applied to.