CONJURA · OPEN PROBLEM

Groth16 Is the Limit in the Pure Generic Group Model

Two lower-bound conjectures on pairing-based SNARGs

Status. Statement: AI-written, not yet formalized.

Category. Proof Systems & Succinct Arguments.

Conjecture (informal). *In the pure generic bilinear group model — no random oracle or any other channel from group elements to bits; the prover can only apply group operations to reference-string elements; the verifier checks pairing product equations fixed by the statement — every sound, nontrivially complete pairing-based SNARG for a hard language has at least three group elements. Groth16 is proof-size optimal.*

1 The setting

In a pairing-based succinct non-interactive argument (SNARG) the prover publishes a few group elements and the verifier checks pairing equations; no shorter proof systems are known. Groth's argument for arithmetic circuit satisfiability [Gro16] — Groth16 — needs only three: two elements of $\mathbb{G}_1$ and one of $\mathbb{G}_2$, verified by a single pairing product equation, knowledge-sound against generic adversaries. The same paper rules out single-element proofs and closes with the problem taken up here: whether the bound extends to two elements, "which would prove optimality of our 3 element construction" [Gro16]. Section 2 formalizes the statement above: in the *purist* reading of the model, where Groth16's own security proof lives, the answer is negative.

Notation. $(p, \mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T, e, [1]_1, [1]_2)$ is an asymmetric (Type III) bilinear group of prime order p: no efficient homomorphism exists in either direction between $\mathbb{G}_1$ and $\mathbb{G}_2$. Brackets denote discrete logarithms: $[a]_i \in \mathbb{G}_i$ for $a \in \mathbb{F} := \mathbb{Z}_p$, groups are written additively, $[a]_1 \cdot [b]_2 = [ab]_T$, and linear algebra lifts entrywise,

$M[a]_i = [Ma]_i$; $\langle \cdot, \cdot \rangle$ is the inner product and $(a; b)$ vertical concatenation. Everything below is publicly verifiable; nothing assumes zero knowledge — impossibilities are for plain soundness, a fortiori for stronger notions.

The purist model. In Shoup's generic group model [Sho97] elements are random strings an adversary may compare, hash, or take apart bit by bit; in Maurer's [Mau05] they live behind handles, touchable only through the group oracles — no channel from group elements back to bits. Groth's construction and lower bounds [Gro16] concern provers that pick matrices and apply them to the CRS "in the exponent", and verifiers that check pairing product equations with proof-independent coefficients. We call such a scheme *purist* when the prover is generic in Maurer's sense with no auxiliary oracles (the name *pure GGM* follows [ADY25]): in particular, no random oracle. Groth flags the CRS-side analogue (a reference string G^b encoding bits b) [Gro16]; fact (F3b) shows that restoring the prover-side channel collapses the proof to two elements. Definition 4 formalizes the model.

What is known. Four facts fix the landscape.

(F1) *Three elements suffice.* Groth16 is, at its core, a split noninteractive linear proof (Definition 2) of dimension $(2, 1)$ for quadratic arithmetic programs: perfectly complete, perfectly zero-knowledge, statistically knowledge-sound against affine provers (Definition 3) [Gro16], in the formalism of [BCIOP13]. Compiled: $2\,\mathbb{G}_1 + 1\,\mathbb{G}_2$ (1536 bits over BLS12-381), one pairing product equation.

(F2) *One element is impossible — robustly.* NILPs with a degree-1 decision procedure do not exist for relation generators with hard decisional problems (Definition 1), so a purist proof needs elements of both source groups, $k_1, k_2 \geq 1$: at least two elements [Gro16, Thms. 3–4]. Assuming one-way functions, Arnon, Dujmovic and Yogev [ADY25] extend the one-element impossibility to *arbitrary* Maurer-generic verifiers (restriction (iii) of Definition 4 dropped) and, for non-adaptive verifier queries and oracle-free setups, to the GGM + ROM.

(F3) *Every purist restriction is load-bearing.* (a) *Symmetry.* In Type I groups ($\mathbb{G}_1 = \mathbb{G}_2$) two elements suffice: squaring gates force

$u_i = v_i$, so $B = A + \beta - \alpha$ and the verifier recomputes $[B]$ itself, leaving (A, C) [Gro16, Sect. 3.1]. Two *field* elements likewise suffice at the bare NILP level, so no dimension count over $\mathbb{F}$ can prove the conjectures: the split is the crux. (b) *Random oracles.* In the GGM + ROM — Maurer's model plus an oracle taking group elements to bits — two elements suffice, both in $\mathbb{G}_1$ (768 bits over BLS12-381): the prover of [ADY25] feeds its first proof element through the oracle to derive the verifier's otherwise-unpredictable query — exactly the channel the purist model lacks — and the underlying PCP has accepting set of size 1, so no answer is sent. (Completeness becomes imperfect; negligible errors need $\lambda \leq \log^c n$, or Khot's Unique Games Conjecture [Kho02], used for completeness, not soundness.) (c) *Non-group data.* Field elements in the proof buy fewer bits with more elements: Polymath, $3\,\mathbb{G}_1 + 1\,\mathbb{F} \approx 1408$ bits [Lip24]; Pari, $2\,\mathbb{G}_1 + 2\,\mathbb{F} \approx 1280$ bits [DMS24]. Element count and bit count thus part ways; the conjectures concern the former. (d) *Interaction; designated verifiers.* One round of interaction yields one-element proofs [BIOW20]; designated-verifier SNARGs reach one element plus $O(\lambda)$ bits [ADI25]. Non-interactivity and public verifiability are load-bearing too.

(F_4) *Simulation extractability is settled.* Pairing-based simulation-extractable SNARKs need at least 3 elements and 2 verification equations, matched by [GM17]; Groth16 proofs are perfectly rerandomizable, so simulation extractability is strictly stronger and the bound does not transfer.

The gap. By F2 a two-element purist proof is exactly $\pi = ([A]_1, [B]_2)$, and in the exponent each verification equation reads

$$c_i \cdot AB \;+\; A\,\langle \boldsymbol{u}_i, \boldsymbol{\sigma}_2 \rangle \;+\; \langle \boldsymbol{v}_i, \boldsymbol{\sigma}_1 \rangle\, B \;+\; \boldsymbol{\sigma}_1^{\top} M_i\, \boldsymbol{\sigma}_2 \;=\; 0.$$

Every term is affine in the prover's coefficients except the single bilinear monomial AB — and the sampling attacks behind [Gro16, Thms. 3–4] and the one-element bound of [ADY25] die on exactly that monomial (Remark 1). The no-ROM two-element question is [ADY25]'s first open problem; the conjectures assert the negative answer.

2 The conjectures

Notation as in Section 1; λ is the security parameter. Definitions follow [Gro16], restated for self-containment; the conjectures quantify over all relation generators.

Definition 1 (relation generator; hard decisional problems [Gro16, Def. 5]). A *relation generator* $\mathcal{R}$ returns on 1^λ a polynomial-time decidable binary relation R (with statements ϕ, witnesses w, and language $L_R := \{\phi : \exists w,\ (\phi, w) \in R\}$) together with auxiliary input z. It has *hard decisional problems* if there are polynomial-time samplers with $\mathsf{Yes}(R) \to (\phi, w) \in R$ and $\mathsf{No}(R) \to \phi \notin L_R$ (with overwhelming probability), such that for all polynomial-time distinguishers $\mathcal{A}$, over $\phi_0 \leftarrow \mathsf{No}(R)$, $(\phi_1, w_1) \leftarrow \mathsf{Yes}(R)$ and $b \leftarrow \{0,1\}$,

$$\Pr\big[\mathcal{A}(R, z, \phi_b) = b\big] \;\approx\; \tfrac{1}{2}.$$

Any NP-complete relation qualifies if one-way functions exist [Gro16].

Definition 2 (split NILP of proof dimension (k_1, k_2) [Gro16, Sect. 2.5]). A *split non-interactive linear proof* for $\mathcal{R}$ is a triple $(\mathsf{Setup}, \mathsf{Prove}, \mathsf{Vfy})$ of the following form. $\mathsf{Setup}(R)$ outputs $\boldsymbol{\sigma} = (\boldsymbol{\sigma}_1, \boldsymbol{\sigma}_2) \in \mathbb{F}^{m_1} \times \mathbb{F}^{m_2}$, each part containing 1 as an entry, and a simulation trapdoor $\boldsymbol{\tau} \in \mathbb{F}^n$. $\mathsf{Prove}(R, \boldsymbol{\sigma}, \phi, w)$ samples $(\Pi_1, \Pi_2) \leftarrow \mathsf{ProofMatrix}(R, \phi, w)$ with $\Pi_j \in \mathbb{F}^{k_j \times m_j}$ and outputs $\boldsymbol{\pi} = (\boldsymbol{\pi}_1, \boldsymbol{\pi}_2) = (\Pi_1 \boldsymbol{\sigma}_1,\ \Pi_2 \boldsymbol{\sigma}_2)$. $\mathsf{Vfy}(R, \boldsymbol{\sigma}, \phi, \boldsymbol{\pi})$ derives deterministic tests $(T_1, \ldots, T_\eta) \leftarrow \mathsf{Test}(R, \phi)$ with $T_i \in \mathbb{F}^{(m_1+k_1)\times(m_2+k_2)}$, and accepts if and only if

$$(\boldsymbol{\sigma}_1; \boldsymbol{\pi}_1)^\top\, T_i\, (\boldsymbol{\sigma}_2; \boldsymbol{\pi}_2) \;=\; 0 \quad \text{for all } i = 1, \ldots, \eta.$$

The NILP has *completeness error* σ if honest proofs for $(\phi, w) \in R$ verify with probability at least $1 - \sigma$; Groth16's NILP is perfectly complete, $\sigma = 0$.

Definition 3 (statistical soundness against affine prover strategies [Gro16, Def. 6]). An *affine prover strategy* chooses (ϕ, Π_1, Π_2) given only (R, z) — in particular, independently of $\boldsymbol{\sigma}$. A split NILP is *statistically sound against affine prover strategies* if for all (unbounded) adversaries $\mathcal{A}$, over $(R, z) \leftarrow \mathcal{R}(1^\lambda)$, $(\boldsymbol{\sigma}, \boldsymbol{\tau}) \leftarrow \mathsf{Setup}(R)$ and $(\phi, \Pi_1, \Pi_2) \leftarrow \mathcal{A}(R, z)$,

$$\Pr\Big[\phi \notin L_R \wedge \mathsf{Vfy}\big(R, \boldsymbol{\sigma}, \phi, (\Pi_1\boldsymbol{\sigma}_1, \Pi_2\boldsymbol{\sigma}_2)\big) = 1\Big] \approx 0.$$

Definition 4 (purist pairing-based argument). A *purist* pairing-based non-interactive argument for $\mathcal{R}$ is a publicly verifiable argument system in which (i) the reference string is a group description plus group elements only, $\sigma = ([\boldsymbol{\sigma}_1]_1, [\boldsymbol{\sigma}_2]_2, [\boldsymbol{\sigma}_T]_T)$; (ii) the prover is generic in Maurer's sense [Mau05] — it touches group elements only through the group, pairing and equality oracles, and never sees an encoding — so its proof is

$$\pi = \big([\Pi_1\boldsymbol{\sigma}_1]_1,\ [\Pi_2\boldsymbol{\sigma}_2]_2,\ [\Pi_T\boldsymbol{\sigma}_T]_T\big)$$

with matrices depending only on R, ϕ and the observed equality pattern; there are no auxiliary oracles: no random oracle, no channel from group elements to bits; and (iii) from ϕ alone the verifier derives matrices T_i and vectors $\boldsymbol{t}_i$ — pairing product equations with proof-independent coefficients — and accepts if and only if

$$[\boldsymbol{\sigma}_1; \boldsymbol{\pi}_1]_1^\top\, T_i\, [\boldsymbol{\sigma}_2; \boldsymbol{\pi}_2]_2 = [\boldsymbol{t}_i \cdot (\boldsymbol{\sigma}_T; \boldsymbol{\pi}_T)]_T \quad \text{for all } i.$$

The *proof size* is $k_1 + k_2 + k_T$; completeness error is as for split NILPs; *soundness*: every polynomial-query purist prover outputs an accepted (ϕ, π) with $\phi \notin L_R$ with negligible probability. Groth16 is a sound, perfectly complete purist argument of size $(2, 1, 0)$ [Gro16].

Conjecture 1 (no two-element split NILPs). *There are no split NILPs with proof dimension $k_1 = k_2 = 1$ and completeness error at most $1 - 1/\mathrm{poly}(\lambda)$ for relation generators with hard deci-*

sional problems. Concretely: any such candidate (Definition 2) yields either an affine prover strategy breaking soundness (Definition 3), or a polynomial-time algorithm distinguishing the Yes- *and* No-*samplers of Definition 1 with noticeable advantage.*

Conjecture 2 (Groth16 optimality in the pure GGM). *Every sound purist pairing-based non-interactive argument (Definition 4) with completeness error at most* $1-1/\mathrm{poly}(\lambda)$ *for a relation generator with hard decisional problems has proof size at least three. Equivalently: two-group-element purist SNARGs do not exist, and Groth16 is proof-size optimal in the pure GGM.*

An affine strategy is a purist prover making no equality tests, so Conjecture 1 implies Conjecture 2; the converse could fail only via reference strings that are not disclosure-free [Gro16, Def. 4, Lemma 1]. A (1,1) split NILP for a hard relation generator refutes Conjecture 1 — and Conjecture 2 if disclosure-free; no zero knowledge is required. Dropping (iii) as well — arbitrary Maurer-generic verifiers, as in [ADY25] — gives a stronger form we also conjecture.

Remark 1 (why the model should have no room for two). Groth16's third element C absorbs the cross terms $As + rB - rs\delta$ from randomizing A and B, while $\alpha\beta$, γ and δ separate statement-from witness-consistency in one equation. The ROM route (F3b) is hide-then-reveal — commit to the first element *before* the test is drawn — while the purist T_i are fixed by ϕ. The one-element bounds run on uniqueness: a proof-dependent zero-test is affine in the lone element and pins a single accepting value, which yes-instance sampling learns [Gro16, ADY25]; with $k_1 = k_2 = 1$ it pins only a Möbius curve of pairs (A, B), which the sampling attack no longer isolates. The conjecture: AB still cannot both couple the sides' witness dependence and resist sampling.

Remark 2 (a half-remembered lead, recorded with low confidence). The author recalls, but cannot reconstruct, a sketch by which a two-element scheme would contradict a complexity conjecture — possibly about *unique witnesses*, à la Valiant–Vazirani [VV86]; a lead, not a claim. Two facts orbit it: the all-parameter construction of [ADY25] assumes the Unique Games Conjecture [Kho02] (a

"unique"-flavoured conjecture on the *constructive* side; the memory may be that, reversed), and by Remark 1 a (1,1) scheme flirts with unique accepting proofs; making this precise — or refuting it — is part of the problem.

3 Bibliography

[ADI25] Gal Arnon, Jesko Dujmovic and Yuval Ishai. *Designated-verifier SNARGs with one group element.* In CRYPTO 2025, pages 192–224. Springer, 2025.

[ADY25] Gal Arnon, Jesko Dujmovic and Eylon Yogev. *Pairing-based SNARGs with two group elements.* Cryptology ePrint Archive, Paper 2025/2160, November 2025 (revised February 12, 2026). https://eprint.iacr.org/2025/2160

[BCIOP13] Nir Bitansky, Alessandro Chiesa, Yuval Ishai, Rafail Ostrovsky and Omer Paneth. *Succinct non-interactive arguments via linear interactive proofs.* In TCC 2013, volume 7785 of LNCS, pages 315–333. Springer, 2013.

[BIOW20] Ohad Barta, Yuval Ishai, Rafail Ostrovsky and David J. Wu. *On succinct arguments and witness encryption from groups.* In CRYPTO 2020, pages 776–806. Springer, 2020.

[DMS24] Michel Dellepere, Pratyush Mishra and Alireza Shirzad. *Garuda and Pari: faster and smaller SNARKs via equifficient polynomial commitments.* Cryptology ePrint Archive, Paper 2024/1245, 2024. https://eprint.iacr.org/2024/1245

[GM17] Jens Groth and Mary Maller. *Snarky signatures: minimal signatures of knowledge from simulation-extractable SNARKs.* In CRYPTO 2017. Springer, 2017. Also: Cryptology ePrint Archive, Paper 2017/540.

[Gro16] Jens Groth. *On the size of pairing-based non-interactive arguments.* In EUROCRYPT 2016, pages 305–326. Springer, 2016. https://doi.org/10.1007/978-3-662-49896-5_11; also: Cryptology ePrint Archive, Paper 2016/260.

[Kho02] Subhash Khot. *On the power of unique 2-prover 1-round games.* In STOC 2002. ACM, 2002.

[Lip24] Helger Lipmaa. *Polymath: Groth16 is not the limit.* In CRYPTO 2024, volume 14929 of LNCS, pages 170–206. Springer, 2024. Also: Cryptology ePrint Archive, Paper 2024/916.

[Mau05] Ueli Maurer. *Abstract models of computation in cryptography.* In Cryptography and Coding (IMACC 2005), volume 3796 of LNCS, pages 1–12. Springer, 2005.

[Sho97] Victor Shoup. *Lower bounds for discrete logarithms and related problems.* In EUROCRYPT 1997, volume 1233 of LNCS, pages 256–266. Springer, 1997.

[VV86] Leslie G. Valiant and Vijay V. Vazirani. *NP is as easy as detecting unique solutions.* Theoretical Computer Science, 47:85–93, 1986.