Non-Interactive Quantum Zero-Knowledge: Scope, Limits, and the Cost of Shared Entanglement in Authentication Protocols

Éricles Antônio Aquiles Barbosa Lima[1]

[1]IMECC – Unicamp
Campinas – SP – Brazil

ericles.aquiles@ime.unicamp.br

Abstract. *Non-interactive quantum statistical zero-knowledge (NIQSZK), introduced by Kobayashi in 2002 via shared EPR pairs replacing a classical common reference string, is often invoked as a general tool for quantum-native authentication, but the term spans results of very different scope. We give an account of what Kobayashi's model actually establishes – non-interactive proofs for one promise problem, Quantum State Closeness to Identity, and, by reduction, Graph Non-Automorphism – and contrast it with the independent, more recent line achieving non-interactive zero-knowledge arguments for all of QMA from quantum-hardness of Learning With Errors (LWE). Using a Software-Defined Quantum Network (SDQN) control-plane authentication case study, we compare the realistically available proof systems along a resource-cost model grounded in published entanglement-distribution benchmarks and lattice-parameter estimation, and argue classical post-quantum NIZK dominates unless the credential itself is an inherently quantum object.*

Resumo. *A prova de conhecimento zero quântica não interativa (NIQSZK), introduzida por Kobayashi em 2002 usando pares EPR no lugar de uma string de referência comum clássica, é frequentemente invocada como ferramenta geral para autenticação nativamente quântica, mas o termo abrange resultados de escopo muito distinto. Apresentamos um relato preciso do que o modelo de Kobayashi de fato estabelece – provas não interativas para um problema de promessa específico, Quantum State Closeness to Identity, e, por redução, Graph Non-Automorphism – contrastando-o com a linha independente e mais recente que obtém argumentos de conhecimento zero não interativos para toda a classe QMA a partir da dificuldade quântica do LWE. Usando um estudo de caso de autenticação no plano de controle de uma SDQN, comparamos os sistemas de prova disponíveis sob um modelo de custo fundamentado em benchmarks de distribuição de emaranhamento e estimativas de parâmetros reticulados, argumentando que o NIZK pós-quântico clássico domina, salvo quando a credencial é um objeto inerentemente quântico.*

1. Introduction

QKD networks offer information-theoretic security for the *data plane*, but say nothing about the classical *control plane* that manages routing and orchestration in a Software-Defined Quantum Network (SDQN). A compromised or spoofed routing node can inject

false topology updates or mount denial-of-service attacks regardless of how secure the underlying quantum channel is, so control-plane authentication is a prerequisite for a quantum network's end-to-end security claims, not an afterthought.

The default answer is post-quantum classical signatures, typically lattice-based. It is natural to ask whether a network that already distributes entangled resources might instead authenticate its control plane with genuinely quantum proof techniques. This sits within the older line of research on *non-interactive quantum statistical zero-knowledge* (NIQSZK), dating to Kobayashi's formulation of non-interactive quantum statistical and perfect zero-knowledge using shared EPR pairs in place of a CRS [Kobayashi 2002].

Revisiting this literature for a concrete application surfaces a problem that is easy to miss: "non-interactive quantum statistical zero-knowledge" spans results of very different scope. Kobayashi's results hold for specific promise problems – Quantum State Closeness to Identity and, via reduction, Graph Non-Automorphism – not NP or QMA in general. A separate, more recent line built on quantum-hardness of Learning With Errors (LWE) has achieved non-interactive zero-knowledge argument systems for all of QMA [Broadbent et al. 2016, Bitansky and Shmueli 2020, Coladangelo et al. 2020, Alagic et al. 2020, Bartusek and Jawale 2025]. Treating "shared EPR pairs replace the CRS" as a general recipe for an arbitrary authentication relation conflates these two lines and may lead to protocol proposals that sound plausible but rest on results the literature does not establish.

This paper (1) gives an account of the NIQSZK landscape, distinguishing Kobayashi's shared-entanglement model from the LWE-based QMA literature (Sections 2–3); and (2) uses this picture to evaluate an SDQN control-plane authentication scenario, comparing available proof systems along a resource-cost model grounded in entanglement-distribution benchmarks and lattice-parameter estimation (Section 4). We close with limitations and open problems (Section 5), including what we believe is the most promising way to make Kobayashi's model practically relevant: a QSCI-style reduction for an authentication-relevant relation, which remains open.

2. Background

Throughout, we use three complexity classes informally: **NP** (languages with a classical witness checkable in polynomial time), **QMA** (its quantum analogue, with a polynomial-size quantum witness and bounded-error quantum verification), and **BQP** (languages decidable outright by a polynomial-time quantum algorithm, with no witness at all).

2.1. Classical Non-Interactive Zero-Knowledge

Given a language L, a zero-knowledge proof lets a prover P convince a verifier V that $x \in L$ without revealing more, formalized via a simulator reproducing V's view without the witness. Goldwasser–Micali–Rackoff required multiple interaction rounds; Blum, Feldman, and Micali showed interaction can be eliminated if P, V share a common reference string (CRS) generated once and reused [Blum et al. 1988]. This *non-interactive zero-knowledge* (NIZK) model removes round-trip latency, letting a node present a single self-contained proof.

Subsequent work gives NIZK for all of NP under plain LWE via Fiat–Shamir with correlation-intractable hash functions [Peikert and Shiehian 2022]. This matters for

Section 4: post-quantum-secure, purely classical NIZK is already achievable from lattice assumptions with no quantum channel at all – the baseline our cost comparison needs.

2.2. Non-Interactive Quantum Statistical Zero-Knowledge à la Kobayashi

Kobayashi [Kobayashi 2002] asks what happens when both proof system and shared setup are quantum, defining NIQSZK/NIQPZK by replacing the CRS with EPR pairs shared before the protocol begins. Two results anchor the model. First, some shared resource – randomness or entanglement – is *necessary*: without it the model collapses to BQP, so a protocol is *non-trivial* (in Kobayashi's sense) exactly when it proves membership in a language outside BQP – otherwise the verifier could decide the instance unaided and the proof is superfluous. Second, Kobayashi gives a natural complete promise problem, Quantum State Closeness to Identity (QSCI): decide whether a given circuit's output is close to maximally mixed or bounded away from it; every NIQSZK language reduces to QSCI. To bring evidence that the model is non-trivial in this sense, Graph Non-Automorphism (GNA) – not known to be in BQP – is shown to admit a perfect-completeness NIQPZK proof via reduction to QSCI.

Crucially, this zero-knowledge guarantee holds against an arbitrary, possibly cheating verifier, strengthening Watrous's earlier interactive results, which held only for honest verifiers [Watrous 2002] (Watrous later showed several *interactive* classical protocols are zero-knowledge against general quantum verifiers too [Watrous 2009], by different means).

Scope matters here: Kobayashi's results are stated for specific promise problems, not NP or QMA in general, and this is more than an open question – Menda and Watrous exhibit an oracle separation suggesting Kobayashi's complete problem is unlikely to contain even a subclass of NP [Menda and Watrous 2018], i.e., a general NIZK for NP is unlikely within this shared-entanglement model.

2.3. Subsequent Developments

The field's center of gravity later shifted to QMA, the quantum analogue of NP. Broadbent, Ji, Song, and Watrous gave zero-knowledge *interactive* proofs for all of QMA, assuming a quantum-concealing commitment scheme [Broadbent et al. 2016]. Bitansky and Shmueli gave a constant-round zero-knowledge argument for QMA from quantum-secure LWE and QFHE, via a non-black-box "no-cloning" simulation technique [Bitansky and Shmueli 2020] – still interactive (multiple messages), and using no shared entangled state at all.

Coladangelo, Vidick, and Zhang gave a NIZK *argument* for QMA with preprocessing: a classical CRS plus one instance-independent quantum message from verifier to prover, reused across proofs, after which verification is classical [Coladangelo et al. 2020]; their model is designated-verifier (secret parameters)[1], and they formalize arguments/proofs of quantum knowledge (AoQK/PoQK). Alagic, Childs, Grilo, and Hung gave the first NIZK for QMA with purely classical communication, secure in the quantum random oracle model under LWE hardness plus circuit-private FHE

[1]In a *designated-verifier* NIZK, only a verifier holding a secret setup parameter can check proofs, unlike *publicly verifiable* NIZK, where anyone can verify using public information alone.

[Alagic et al. 2020]. Most recently, Bartusek and Jawale gave the first publicly verifiable, transparent-setup NIZK argument of quantum knowledge for QMA, with witness extractability [Bartusek and Jawale 2025] – though its security relies on the *heuristic* use of a post-quantum indistinguishability obfuscator (or a heuristically instantiated quantum pseudorandom oracle), not a clean reduction to plain LWE.

Two points carry forward: the field moved *away* from EPR-pair-as-setup and *toward* LWE-based constructions where "quantumness" lives in the witness, not a distributed entangled resource; and none of these constructions inherits Kobayashi's free cheating-verifier guarantee – each makes its own assumptions (designated-verifier, random oracle, or heuristic obfuscation).

3. The Gap Between Promise Problems and General-Purpose Quantum NIZK

Kobayashi's NIQSZK/NIQPZK results are proven for one complete promise problem, QSCI, and one language reducing to it, GNA – not for arbitrary NP or QMA languages. As noted above, Menda and Watrous's oracle separation gives evidence this scope is not accidental: QSCI-completeness is unlikely to extend to (even a subclass of) NP [Menda and Watrous 2018]. So "Kobayashi shows NIQSZK is possible via shared EPR pairs" does not imply an arbitrary authentication relation inherits a non-interactive quantum statistical ZK proof this way; using the EPR-pairs-as-CRS mechanism specifically requires a new QSCI reduction, which does not currently exist in the literature as far as we have found.

The QMA literature (Bitansky–Shmueli, Coladangelo–Vidick–Zhang, Alagic–Childs–Grilo–Hung, Bartusek–Jawale) achieves the generality Kobayashi's paper leaves open, but via a different route: hardness comes from quantum-hardness of LWE (plus, in the newest result, heuristic obfuscation), not promise-problem structure; the quantumness lives in the witness, while communication in the strongest results is classical or uses one consumed-once quantum message; and cheating-verifier security is not free – each construction has its own setup-specific guarantee. In short, for a general-purpose, minimal-quantum-infrastructure NIZK today, the LWE-based QMA literature is the applicable toolkit, not the EPR-shared-state model, despite both sharing the "non-interactive quantum statistical zero-knowledge" label.

This creates two design points, carried into Section 4: (1) stay inside Kobayashi's model, viable only if the authorization relation reduces to a QSCI-like problem – we know of no such reduction for credential possession, and flag it as future work; or (2) use an LWE-based QMA/NP NIZK construction, sacrificing the free cheating-verifier guarantee but available today for general relations, with several strong variants (ACGH, Bartusek–Jawale) needing little or no persistent quantum resource.

4. Case Study: Node Authentication in an SDQN Control Plane

4.1. Scenario and candidate constructions

Consider an SDQN where routing nodes repeatedly authenticate to a controller for topology updates, proving "this node holds a credential issued by a trusted authority" without revealing which one. Repetition – not a single proof – is what makes cost comparisons matter.

(i) Kobayashi-style EPR-based NIQSZK requires reducing the authorization relation to QSCI; we know of no such reduction for credential possession, so we exclude it from the quantitative comparison and list it as future work.

(ii) LWE-based NIZK/argument for QMA (CVZ20-style or its classical-communication/transparent-setup successors) applies since credential-possession is in NP (or QMA, if the credential is a quantum state, e.g. an unclonable token). CVZ20 needs a one-time quantum preprocessing message in a designated-verifier model; ACGH and Bartusek–Jawale remove it, at the cost of a random-oracle or heuristic-obfuscation assumption. All rest primarily on quantum-hardness of LWE.

(iii) Classical post-quantum NIZK (Peikert–Shiehian-style) suffices if the credential is classical, needs no quantum channel, and runs entirely over existing classical control-plane links.

4.2. Resource-cost model

Entanglement setup cost (ii only). A quantum preprocessing message that is not persistent must be regenerated as entangled resources decohere. Metropolitan fiber testbeds have demonstrated continuous distribution around 2×10^4 pairs/s at $\sim$99% fidelity sustained over 15 days, and up to 5×10^5 pairs/s in short high-rate windows on a deployed 34 km fiber [Craddock et al. 2024]. Longer links cost fidelity and rate steeply: 0.7 pairs/s at 155 km versus 132 pairs/s at 93 km on deployed fiber with a silicon photonic source [Du et al. 2024]. Practically, reused CVZ20-style preprocessing is plausible with existing metro-scale hardware, whereas a scheme needing a *fresh* EPR pair per authentication event would need rates/fidelities current deployments only reach at short-to-medium range – favoring the "preprocessing, reused" model over a hypothetical per-proof EPR pair.

Lattice-parameter cost (ii and iii). Both rest on LWE hardness at dimension n, modulus q, noise χ. We defer a full numeric table to an extended version, since parameter sets are construction-specific and not uniformly reported; the point is that this comparison can be made *quantitative* with tooling the lattice-cryptography community already uses (such as [Albrecht et al. 2015]), not left qualitative.

4.3. Summary comparison

Table 1. LWE-based NIZK for QMA vs. classical post-quantum NIZK for NP.

	(ii) LWE-based NIZK for QMA	(iii) Classical PQ-NIZK for NP
Quantum channel	Sometimes (reusable preprocessing in CVZ20; none in ACGH/Bartusek–Jawale)	No
Per-proof entanglement cost	None beyond setup	None
Security basis	LWE (+RO, designated-verifier, or heuristic obfuscation)	LWE
Credential model	Classical or quantum (QMA)	Classical only (NP)
Deployment readiness	Plausible for one-time setup at metro distances	Immediate, existing links

Unless the credential is inherently a quantum object, (iii) dominates for SDQN control-plane authentication: no quantum channel, and its parameters can be analyzed

with the same lattice-estimation tooling as (ii). Option (ii) is preferable specifically when the credential must be a quantum state – a design requirement to justify explicitly, not a default.

5. Discussion

The contribution is (a) a corrected account of what Kobayashi's NIQSZK model does and does not establish, and (b) a resource-grounded argument for which *available* proof system fits a realistic use case – worth stating plainly, since the literature invites reading "quantum zero-knowledge" as one tool when it is several, with different scopes, assumptions, and deployment costs (Section 4).

Limitations. Our cost model is first-order: it omits multi-hop entanglement swapping and fidelity loss, quantum-memory cost if a CVZ20-style message cannot be consumed immediately, and credential-revocation overhead, which none of the surveyed constructions address natively. We also did not complete the BKZ-based numeric table promised in Section 4; this requires a specific published parameter set and is a natural next step. Our exclusion of option (i) reflects an absent reduction (and evidence against a general one [Menda and Watrous 2018]), not a proof of impossibility for every relation.

Open problems. (1) A QSCI-style reduction for a security-relevant relation – GNA remains the only example beyond QSCI itself; graph-isomorphism-based identification schemes may be a promising target, though [Menda and Watrous 2018] suggests any such reduction needs special structure, not a generic NP-hard relation. (2) Clarifying what cheating-verifier guarantees, if any, the CVZ20/ACGH/Bartusek–Jawale constructions provide, relevant when an SDQN controller cannot be fully trusted with prover privacy. (3) Empirically measuring CVZ20-style preprocessing cost end-to-end on a deployed metropolitan testbed, replacing our order-of-magnitude estimate.

For practitioners: unless the credential is itself quantum, classical post-quantum NIZK needs no new infrastructure and can be parameter-tuned with existing lattice tooling; quantum-native proof systems earn their place only when the object being proven cannot be represented classically.

6. Conclusion

Kobayashi's non-interactive quantum statistical zero-knowledge model and the subsequent LWE-based QMA literature are often conflated under one label despite resting on different mechanisms – shared entanglement versus quantum-hardness of LWE – and different scopes, a gap reinforced by an oracle separation suggesting Kobayashi's model does not extend to general NP [Menda and Watrous 2018]. In an SDQN control-plane case study, the practically deployable options today are LWE-based, and the choice between a classical post-quantum NIZK and a QMA-level construction should be driven by whether the credential is a quantum object – a concrete design question, not a default. A reduction extending Kobayashi's model to a security-relevant relation, and an empirical measurement of preprocessing cost on a deployed testbed, remain for future work.

References

Alagic, G., Childs, A. M., Grilo, A. B., and Hung, S.-H. (2020). Non-interactive classical verification of quantum computation. In *Theory of Cryptography Conference*

(TCC 2020), Part III, volume 12552, pages 153–180. Springer. Also Cryptology ePrint Archive, Report 2020/1422; arXiv:1911.08101.

Albrecht, M. R., Player, R., and Scott, S. (2015). On the concrete hardness of learning with errors. *J. Math. Cryptol.*, 9(3):169–203.

Bartusek, J. and Jawale, R. (2025). A new approach to arguments of quantum knowledge. Cryptology ePrint Archive, Report 2025/2155; arXiv:2510.05316.

Bitansky, N. and Shmueli, O. (2020). Post-quantum zero knowledge in constant rounds. In *Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing (STOC)*, pages 269–279.

Blum, M., Feldman, P., and Micali, S. (1988). Non-interactive zero-knowledge and its applications. In *Proceedings of the 20th Annual ACM Symposium on Theory of Computing (STOC)*, pages 103–112.

Broadbent, A., Ji, Z., Song, F., and Watrous, J. (2016). Zero-knowledge proof systems for QMA. In *IEEE 57th Annual Symposium on Foundations of Computer Science (FOCS 2016)*, pages 31–40. Journal version: SIAM Journal on Computing, 49(2):245–283, 2020.

Coladangelo, A., Vidick, T., and Zhang, T. (2020). Non-interactive zero-knowledge arguments for QMA, with preprocessing. In *Advances in Cryptology – CRYPTO 2020*, volume 12172, pages 799–828. Springer.

Craddock, A. N., Lazenby, A., Portmann, G. B., Sekelsky, R., and Namazi, M. (2024). Automated distribution of high-rate, high-fidelity polarization entangled photons using deployed metropolitan fibers. arXiv:2404.08626.

Du, J. et al. (2024). Entanglement distribution over 155 km metropolitan fiber using a CMOS-compatible silicon chip. arXiv:2409.17558.

Kobayashi, H. (2002). Non-interactive quantum statistical and perfect zero-knowledge. arXiv:quant-ph/0207158. Published as "Non-interactive Quantum Perfect and Statistical Zero-Knowledge," in Algorithms and Computation, 14th International Symposium (ISAAC 2003), LNCS vol. 2906, pp. 178–188, Springer, 2003.

Menda, S. and Watrous, J. (2018). Oracle separations for quantum statistical zero-knowledge. arXiv:1801.08967.

Peikert, C. and Shiehian, S. (2022). Noninteractive zero knowledge for NP from (plain) learning with errors. In *Advances in Cryptology – CRYPTO 2022*. Springer.

Watrous, J. (2002). Limits on the power of quantum statistical zero-knowledge. In *Proceedings of the 43rd Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 459–468.

Watrous, J. (2009). Zero-knowledge against quantum attacks. *SIAM Journal on Computing*, 39(1):25–58. Preliminary version in STOC 2006.