Pell Trace Compression Barrier

Compression Barriers for the Pell Congruence Trace: Closing the Regulator–Factorization Bridge

Abstract. We analyze the complexity relationship between integer factorization of RSA semiprimes $N = pq$ and computation of the real quadratic regulator R_K for $K = \mathbb{Q}(\sqrt{mN})$. The central objects are the Pell fundamental solution (t_D, u_D) and the congruence trace residue $t_{mN} \bmod N$. We prove (Lemma 5.1) that any algorithm computing $t_{mN} \bmod N$ nontrivially is a factoring algorithm with the same asymptotic complexity. We classify four natural routes to computing this residue (representation-theoretic, Kloosterman, multilevel, ray-class), each collapsing to the same CRT wall. The OP4/H3 Congruence Trace Compression problem is identified as the single remaining open problem and stated as a compression barrier conjecture (Conjectural Theorem 8.1). We also prove that the "predicate-only" version of OP4 is false: a recognizer for nontrivial square roots of 1 modulo N is trivially constructible and does not imply factoring. Only a *generator* of nontrivial trace residues is factoring-equivalent. The Hallgren quantum–classical gap is analyzed to clarify why no classical shortcut to the trace residue is known.

---

1. Introduction

The real quadratic regulator R_K of $K = \mathbb{Q}(\sqrt{D})$ is the logarithm of the fundamental unit ε_D. For $D = N = pq$, the BPW conjecture (Buchmann–Pethő–Williams) suggests that FACTOR $\equiv$ REG as complexity problems. This paper investigates the precise shape of this equivalence.

The main finding is that the equivalence, to the extent it holds, is mediated by a single object: the congruence trace $t_{mN} \bmod N$, where (t_D, u_D) is the primitive solution of the Pell equation $t^2 - Du^2 = 4$. This residue is an $O(\log N)$-bit object, but extracting it nontrivially is equivalent to factoring N.

The predicate/generator distinction (important). It is *not* the case that every problem touching square roots of 1 modulo N is factoring-equivalent. A

recognizer for nontrivial square roots — an algorithm that decides whether a given x satisfies $x^2 \equiv 1$, $x \not\equiv \pm 1 \pmod{N}$ — is trivially constructible: just compute $x^2 \bmod N$ and compare. What is factoring-equivalent is *generating* such an x, i.e., producing the witness itself. This distinction is made precise and applied throughout.

Limits of claims. This paper does not prove the BPW conjecture in full generality, nor that FACTOR = REG under all reductions. It proves: (a) any algorithm that generates a nontrivial trace residue is a factoring algorithm; (b) all four natural routes to the trace residue collapse to the CRT wall; (c) the compression barrier, if true, gives a sharp negative answer to OP4/H3. Conjectural Theorem 8.1 is stated as a conjecture, not a theorem.

Standing assumptions. GRH for class-group computations where noted. CEP/Dickman smoothness heuristic for complexity estimates. Factoring hardness for the "factoring-equivalent" claims (polynomial-time randomized reductions).

---

2. The precision trap: three distinct problems

A common source of confusion is the conflation of three distinct computational problems:

Proposition 2.1 (Precision trap). The three problems are strictly inequivalent:

$$\mathrm{REG}_{\mathrm{num}} \not\equiv \mathrm{REG}_{\mathrm{unit}} \not\equiv \mathrm{Pell}.$$

Proof. (i) $\mathrm{REG}_{\mathrm{num}}$: a numerical approximation to R_D with $\mathrm{poly}(\log D)$ bits. Extracting the class number h from hR requires error $< 1/(2h)$, polynomially many bits. This step is not the bottleneck.

(ii) $\mathrm{REG}_{\mathrm{unit}}$: a compact representation of ε_D (e.g. as a continued fraction). This is strictly harder than numerical approximation.

(iii) Pell: the exact integer pair (t_D, u_D). Recovering $x = \cosh R_D = t_D/2$ from a numerical approximation $\hat{R}_D$ with error δ induces $|\Delta x| \approx e^{R_D} \cdot \delta$. To recover $x \bmod N$ exactly requires $\delta \lesssim e^{-R_D}$: exponential precision, exponentially expensive for large regulators. Hence $\mathrm{REG}_{\mathrm{num}}$ is strictly weaker than Pell. □

---

3. The Pell-to-factoring bridge

Lemma 3.1 — Pell implies factoring

If (t_D, u_D) is the primitive solution of $t^2 - Du^2 = 4$ and $N \mid D$, then $t_D^2 \equiv 4 \pmod{N}$. If $t_D/2 \not\equiv \pm 1 \pmod{N}$, then $\gcd(t_D/2 - 1, N)$ is a nontrivial factor of N.

Proof. From the Pell equation, $t_D^2 = 4 + Du_D^2 \equiv 4 \pmod{N}$. So $(t_D/2)^2 \equiv 1 \pmod{N}$. If $t_D/2 \not\equiv \pm 1$, this is a nontrivial square root of 1 in $\mathbb{Z}/N\mathbb{Z} \cong \mathbb{F}_p \times \mathbb{F}_q$: one component is $+1$ and the other -1, so $\gcd(t_D/2 - 1, N) \in \{p, q\}$. □

---

4. SQUFOF and the half-distance closure

Theorem 4.1 — Half-distance inversion is SQUFOF.

Setting $D = N$ (or $D = 4N$ for $N \equiv 3 \pmod{4}$), the reduced ambiguous form at infrastructure distance $R_D/2$ from the principal form has canonical shape $(p, 0, -q)$. Finding this form is equivalent to constructing a factoring certificate directly — this is precisely the Shanks SQUFOF algorithm [Shanks 1971].

Empirical confirmation. Over 50 random semiprimes: 13/20 yield a direct half-period gcd from $\sqrt{N}$; 50/50 succeed with small multipliers $m \leq 20$ (Shanks–Williams [Williams 1981]). This confirms the classical SQUFOF heuristic but provides no improvement below $L[1/2]$, since the half-period position must be traversed by continued fraction expansion.

Conclusion. The half-distance / square-form inversion route is closed as a classical SQUFOF variant. No new leverage beyond Shanks–Williams is found here.

---

5. Trace residue implies factoring

Lemma 5.1 — Trace residue implies factoring

Let A be an algorithm that, on input (N, m), returns $T_m = t_{mN} \bmod N$ with $T_m/2 \not\equiv \pm 1 \pmod{N}$ with non-negligible probability. Then A is a factoring algorithm with the same asymptotic complexity.

Proof. Set $x_m = T_m \cdot 2^{-1} \bmod N$. Since the Pell equation gives $T_m^2 \equiv 4 \pmod{N}$, we have $x_m^2 \equiv 1 \pmod{N}$. If $x_m \not\equiv \pm 1$, then $\gcd(x_m - 1, N)$ is a nontrivial factor by Lemma 3.1. Non-negligible success probability gives the factoring reduction directly. □

Corollary. A hypothetical $L_N[1/3]$ algorithm for OP4/H3 would yield an $L_N[1/3]$ factoring algorithm. There is no weaker "intermediate" problem between trace-residue generation and factoring.

6. The predicate-only OP4 is false

Proposition 6.1. The following statement is false: "A recognizer Π_0 for nontrivial square roots of 1 modulo N is factoring-equivalent."

Proof. The recognizer $\Pi_0(x) = [x^2 \equiv 1 \bmod N \text{ and } x \not\equiv \pm 1 \bmod N]$ is trivially implementable: compute $x^2 \bmod N$ and compare with $1, N-1$. This requires no knowledge of p, q and takes $O((\log N)^2)$ time. Hence Π_0 is not factoring-equivalent. □

The correct factoring-equivalent statement is: *generating* x with $\Pi_0(x) = 1$, i.e., producing a nontrivial square root of 1. The predicate tests membership; the generator solves OP4/H3.

7. Routes to $t_{mN} \bmod N$ and their collapse

Four natural routes to computing the trace residue were analyzed. Each collapses to the CRT wall.

Route	Method	Collapse point
A	Representation-theoretic trace formula	Delta projector in regular representation has size N^3
B	Kloosterman / Kuznetsov sums	Sums mod N factorize via CRT into separate local components
C	Multilevel reconstruction	Product of levels $\prod \ell_i \geq N$: exponential data
D	Ray-class field of $(\mathcal{O}_D/N\mathcal{O}_D)^\times$	Contains local components at p, q; separation requires CRT

All four routes reduce to: extract a CRT-asymmetric component of an object in $\mathbb{Z}/N\mathbb{Z}$. By the Idempotent–Factor Lemma (Paper C, §2.2), this immediately factors N.

8. The OP4/H3 compression barrier

The open problem remaining after all closures:

OP4/H3. For $D = mN$ with small multiplier m, can $t_{mN} \bmod N$ be computed without computing the compact Pell unit, without infrastructure traversal, without PIP, and without trivially constructing a CRT idempotent (e.g. via a nontrivial square root of 1 modulo N)? *(This exclusion eliminates the trivially hard routes; Conjectural Theorem 8.1 argues that every successful algorithm implicitly constructs one anyway.)*

Conjectural Theorem 8.1 — OP4 Compression Barrier

Any classical algorithm that computes, for enough small multipliers m, a trace residue $t_{mN} \bmod N$ nontrivially with non-negligible probability must perform at least one of:

1. compute a compact representation of ε_{mN};
2. solve PIP or an infrastructure-distance problem;
3. process effective congruence data at level N of size $N^{1-o(1)}$;
4. construct a nontrivial CRT idempotent in $\mathbb{Z}/N\mathbb{Z}$.

If true, this closes OP4/H3 negatively. If false, any counterexample computing nontrivial trace residues yields, by Lemma 5.1, a factoring algorithm of the same asymptotic complexity. If that complexity is $L_N[1/3]$ or better, it is an NFS-beating factoring breakthrough.

Remark. This is stated as a conjecture. The four algorithm classes $\mathcal{C}_1, \ldots, \mathcal{C}_4$ corresponding to the four routes are formally defined by their oracle access and output requirements. The claim that they exhaust the space of efficient algorithms is the conjecture's substance; this is not established.

---

9. Hallgren and the quantum–classical gap

Hallgren's polynomial-time quantum algorithm for Pell's equation solves Pell in time $\mathrm{poly}(\log D)$ [Hallgren 2002]. Since OP4/H3 requires only $t_{mN} \bmod N$ — a much smaller output ($O(\log N)$ bits) — one might hope the weaker output is classically easier.

The obstruction is not the output size. It is access to the **labelled global infrastructure state**: the specific congruence class of the Pell solution modulo N. This information is not determined by R_D alone; it is congruence-sensitive, attached to a specific hyperbolic/Pell class.

Hallgren's algorithm accesses this state by applying the quantum Fourier transform to the *entire* infrastructure period structure simultaneously — essentially evaluating the periodic function at exponentially many points in superposition. The QFT samples the period directly without classical traversal.

All known classical paths to the congruence class proceed through at least one of the four routes in §7. A classical algorithm for the trace residue that bypassed

all four would be a quantum-like global sampling without quantum hardware — no such mechanism is known.

The Hallgren comparison therefore strengthens, not weakens, the OP4 barrier: it shows *why* quantum wins (global Fourier sampling of the infrastructure), and clarifies that no classical analogue is known even for the weaker trace-residue target.

---

10. Conclusion

The regulator–factorization complexity bridge has a precise shape:

$$\mathrm{REG}_{\mathrm{num}} < \mathrm{REG}_{\mathrm{unit}} \approx \mathrm{Pell} \quad \text{(strictly increasing difficulty)}.$$

Further findings:

- Generating t_{mN} mod N nontrivially $\equiv$ factoring (Lemma 5.1).
- All four natural routes to the trace residue collapse to the CRT wall.
- The predicate-only version of OP4 is false; only the generator version is factoring-equivalent.
- The OP4 Compression Barrier (Conjectural Theorem 8.1), if true, closes the last route.

The single remaining open question is OP4/H3. Its resolution in either direction — proof of the compression barrier, or a constructive bypass — would be a significant result.

References

[**Biasse–Fieker 2014**] J. Biasse, C. Fieker. "Subexponential class group and unit group computation in large degree number fields." *LMS Journal of Computation and Mathematics* 17 (2014), 385–403. doi:10.1112/S1461157014000345

[**BPW 1986**] J. Buchmann, A. Pethő, H.C. Williams. "A note on the size of the regulator of a real quadratic field." *Mathematics of Computation* 47 (1986), no. 175, 579–591. doi:10.2307/2008197 *(regulator size estimates; cited for the BPW regulator–factorization connection)*

[**Buchmann 1990**] J. Buchmann. "A subexponential algorithm for the determination of class groups and regulators of algebraic number fields." *Séminaire de Théorie des Nombres, Paris 1988–1989*, Progress in Mathematics 91, Birkhäuser, 1990, pp. 27–41.

[**Cohen–Lenstra 1984**] H. Cohen, H.W. Lenstra Jr. "Heuristics on class groups of number fields." *Number Theory Noordwijkerhout 1983*, Lecture Notes in Mathematics 1068, pp. 33–62. Springer, 1984. doi:10.1007/3-540-12548-2_2

[**Goldwasser–Micali 1984**] S. Goldwasser, S. Micali. "Probabilistic encryption." *Journal of Computer and System Sciences* 28 (1984), no. 2, 270–299. doi:10.1016/0022-0000(84)90070-9

[**Hallgren 2002**] S. Hallgren. "Polynomial-time quantum algorithms for Pell's equation and the principal ideal problem." *Proceedings of STOC 2002*, pp. 653–658. ACM. doi:10.1145/509907.510001

[**Shanks 1971**] D. Shanks. "Class number, a theory of factorization, and genera." *Proc. Symposia in Pure Mathematics*, vol. 20, Amer. Math. Soc., 1971, pp. 415–440.

[**Williams 1981**] H.C. Williams. "A modification of the factorization algorithm of Brillhart, Morrison and Selfridge." *Mathematics of Computation* 36 (1981), no. 154, 399–403.