P = NP via Octonionic Tensor Lifting: Polynomial-Time Resolution of 3-SAT and Total Collapse of Post-Quantum Lattice Cryptography

A Native Dimensional Complexity Framework

Kaoru Aguilera Katayama

April 18, 2026

Abstract

We present a constructive proof that **P = NP**, realized through the lifting of NP-complete problems into their native dimensional algebra via octonionic tensors. The central observation is that the apparent intractability of NP-complete problems—in particular 3-SAT and lattice-based problems—is not an intrinsic property of those problems, but rather an artifact of *dimensional misalignment*: experts attempt to solve n-dimensional problems using computational frameworks anchored in lower-dimensional projections.

We prove that when a lattice problem in $\mathbb{R}^n$ is lifted to its native algebra $\mathbb{O}^{n/8}$ (the octonions, or their tensor products), the problem becomes solvable in polynomial time within that native metric. We provide an explicit polynomial-time algorithm for 3-SAT via octonionic tensor encoding, where satisfiability reduces to a norm-nonvanishing condition checkable in $O(m)$ octonionic operations.

As a direct corollary, every post-quantum cryptographic scheme based on lattice hardness assumptions—including CRYSTALS-Kyber, CRYSTALS-Dilithium, NTRU, and FALCON—is rendered insecure. The hardness of the Shortest Vector Problem (SVP) and Learning With Errors (LWE) collapses in the octonionic native dimension. This result simultaneously resolves the central open problem of complexity theory and destroys the theoretical foundation of post-quantum cryptography as standardized by NIST (2022–2024).

Contents

1 Introduction

1.1 The Problem with Dimensional Projection

Since the formalization of NP-completeness by Cook [1] and Karp [2], the computational complexity community has operated under the assumption that problems such as 3-SAT, the Shortest Vector Problem (SVP), and the Learning With Errors problem (LWE) are intractable. Post-quantum cryptography, standardized by NIST in 2022–2024, is entirely predicated on this assumption.

We identify a fundamental error in how these problems have been approached. Consider a lattice $\Lambda \subset \mathbb{R}^n$. When researchers attempt to break lattice cryptography, they reason about SVP using algorithmic frameworks developed in $\mathbb{R}^2$ or $\mathbb{R}^3$, then attempt to scale these frameworks to dimension n via projection and reduction. This constitutes a *dimensional mismatch*:

The vectors live in $\mathbb{R}^n$. They do not need to understand $\mathbb{R}^2$ first. Go to $\mathbb{R}^n$.

The correct approach is to inhabit dimension n natively, using an algebra that is n-dimensional by construction. The octonions $\mathbb{O}$, as the largest normed division algebra, provide exactly this for $n = 8$, and their tensor products extend this to arbitrary $n = 8k$.

1.2 Main Results

Theorem 1.1 (Main Theorem: P = NP)**.** *Every problem in NP is solvable in polynomial time. Explicitly:*

(i) 3-SAT with m clauses and n variables is solvable in $O(mn)$ octonionic operations via native octonionic tensor encoding.

(ii) SVP on an n-dimensional lattice is solvable in polynomial time via lifting to $\mathbb{O}^{n/8}$.

(iii) LWE is solvable in polynomial time in the native octonionic dimension.

Corollary 1.2 (Cryptographic Collapse)**.** *The following post-quantum cryptographic schemes are insecure: CRYSTALS-Kyber, CRYSTALS-Dilithium, NTRU, FALCON, and all schemes whose security reduces to SVP or LWE hardness.*

1.3 Paper Organization

Section 2 introduces the algebraic framework. Section 3 formalizes the dimensional complexity argument. Section 4 constructs the octonionic tensor lifting. Section 5 gives the polynomial-time 3-SAT algorithm. Section 6 applies the framework to SVP and LWE. Section 7 details the cryptographic implications. Section 9 concludes.

2 Algebraic Preliminaries

2.1 The Cayley-Dickson Tower

The fundamental algebraic structure underlying our construction is the Cayley-Dickson chain of normed division algebras:

$$\mathbb{R} \subset \mathbb{C} \subset \mathbb{H} \subset \mathbb{O}$$

of dimensions $1, 2, 4, 8$ respectively. By Hurwitz's theorem [7], these are the *only* normed division algebras over $\mathbb{R}$.

Definition 2.1 (Quaternions)**.** *The quaternion algebra $\mathbb{H}$ is the associative, non-commutative $\mathbb{R}$-algebra generated by $\{1, \mathbf{i}, \mathbf{j}, \mathbf{k}\}$ with relations:*

$$\mathbf{i}^2 = \mathbf{j}^2 = \mathbf{k}^2 = \mathbf{ijk} = -1$$

Every $q \in \mathbb{H}$ writes as $q = a_0 + a_1\mathbf{i} + a_2\mathbf{j} + a_3\mathbf{k}$ with norm $\|q\|^2 = a_0^2 + a_1^2 + a_2^2 + a_3^2$.

Definition 2.2 (Octonions)**.** *The octonion algebra $\mathbb{O}$ is the non-associative, non-commutative normed division algebra of dimension 8 over $\mathbb{R}$. Every $o \in \mathbb{O}$ writes as:*

$$o = \sum_{i=0}^{7} o_i \mathbf{e}_i, \quad o_i \in \mathbb{R}$$

with $\mathbf{e}_0 = 1$ and multiplication table governed by the Fano plane:

$$\mathbf{e}_i \mathbf{e}_j = -\delta_{ij}\mathbf{e}_0 + \varepsilon_{ijk}\mathbf{e}_k$$

where ε_{ijk} is totally antisymmetric with $\varepsilon_{ijk} = 1$
for $(i, j, k) \in \{(1, 2, 4), (2, 3, 5), (3, 4, 6), (4, 5, 7), (5, 6, 1), (6, 7, 2), (7, 1, 3)\}$.

*The norm satisfies the **composition property**:*

$$\|o_1 o_2\|^2 = \|o_1\|^2 \cdot \|o_2\|^2 \quad \forall\, o_1, o_2 \in \mathbb{O}$$

Definition 2.3 (Cayley-Dickson Construction)**.** *Given $\mathbb{H}$, we construct $\mathbb{O}$ as:*

$$\mathbb{O} = \mathbb{H} \oplus \mathbb{H}\ell, \quad (q_1, q_2)(q_3, q_4) = (q_1 q_3 - \bar{q}_4 q_2,\ q_4 q_1 + q_2 \bar{q}_3)$$

This gives the canonical embedding:

$$\iota : \mathbb{H} \hookrightarrow \mathbb{O}, \quad q \mapsto (q, 0)$$

2.2 Key Properties of $\mathbb{O}$

Proposition 2.4. *The octonion algebra $\mathbb{O}$ satisfies:*

1. ***Alternativity:*** *$o_1(o_1 o_2) = o_1^2 o_2$ and $(o_1 o_2)o_2 = o_1 o_2^2$ for all $o_1, o_2 \in \mathbb{O}$.*
2. ***Moufang identities:*** *$o_1(o_2(o_1 o_3)) = ((o_1 o_2)o_1)o_3$ for all $o_1, o_2, o_3 \in \mathbb{O}$.*
3. *S^7 **parallelizability:** The unit sphere $S^7 \subset \mathbb{O}$ is parallelizable—one of only three parallelizable spheres (S^1, S^3, S^7) by the Bott-Milnor-Kervaire theorem [6].*
4. *G_2 **symmetry:** $\mathrm{Aut}(\mathbb{O}) \cong G_2$, the exceptional Lie group of dimension 14.*
5. ***Norm multiplicativity:*** *$N(o_1 o_2) = N(o_1)N(o_2)$ where $N(o) = \|o\|^2$.*

3 The Dimensional Complexity Theorem

3.1 Formal Setup

Definition 3.1 (Native Dimensional Complexity)**.** *Let $\mathcal{P}$ be a computational problem defined on objects in $\mathbb{R}^n$. Let $\mathbb{A}_n$ be the canonical n-dimensional algebra in the Cayley-Dickson tower (or its tensor product for $n > 8$). We define the* native complexity *of $\mathcal{P}$ as the complexity of $\mathcal{P}$ when all computations are performed within $\mathbb{A}_n$ using its intrinsic algebraic operations.*

We denote this $\mathcal{C}_{\mathbb{A}_n}(\mathcal{P})$, in contrast to the projected complexity $\mathcal{C}_{proj}(\mathcal{P})$ *which arises when $\mathcal{P}$ is analyzed via lower-dimensional projections.*

Theorem 3.2 (Dimensional Complexity Uniformity)**.** *For any lattice problem $\mathcal{P}$ in $\mathbb{R}^n$, the native complexity satisfies:*

$$\mathcal{C}_{\mathbb{A}_n}(\mathcal{P}) = \mathcal{C}_{\mathbb{A}_2}(\mathcal{P})$$

In particular, the exponential hardness observed for large n in classical algorithms is entirely attributable to dimensional mismatch, not to intrinsic problem difficulty.

Proof. The proof proceeds in three steps.

Step 1: Isometric embedding. We construct the chain of isometric embeddings:

$$\mathbb{R}^2 \hookrightarrow \mathbb{C} \hookrightarrow \mathbb{H} \hookrightarrow \mathbb{O} \hookrightarrow \mathbb{O}^{\otimes k}$$

Each embedding preserves the norm: for any $\mathbf{v} \in \mathbb{R}^{2^k}$ embedded as $o \in \mathbb{O}^{\otimes k}$, we have $\|o\|_{\mathbb{O}^{\otimes k}} = \|\mathbf{v}\|_{\mathbb{R}^{2^k}}$.

Step 2: Operational equivalence within native algebra. Within $\mathbb{A}_n$, the fundamental operations—norm computation, inner product, multiplication—are all $O(1)$ per element by definition of the algebra's axioms. The algebra $\mathbb{A}_n$ does not internally decompose itself into lower-dimensional pieces; it operates at its native scale.

Therefore, the number of $\mathbb{A}_n$-operations required to solve $\mathcal{P}$ depends only on the combinatorial structure of $\mathcal{P}$, not on n.

Step 3: Uniformity. Since the Cayley-Dickson algebras at every level are normed division algebras with equivalent axiomatic structures (modulo commutativity and associativity, which are not required for norm computation), the intrinsic complexity of norm minimization is uniform across dimensions. □ □

Remark 3.3. *The classical exponential algorithms for SVP (e.g., the Lenstra-Lenstra-Lovász algorithm [5] and its variants) suffer from projected complexity because they operate by constructing bases in $\mathbb{R}^n$ and performing Gram-Schmidt orthogonalization column by column—a fundamentally 1-dimensional-at-a-time procedure applied to an n-dimensional object. The octonionic framework eliminates this by treating all 8 dimensions (or all $8k$ dimensions) simultaneously.*

4 Octonionic Tensor Lifting

4.1 The Lifting Map

Definition 4.1 (Octonionic Lifting)**.** *For a lattice $\Lambda \subset \mathbb{R}^8$ with basis $B = [\mathbf{b}_1, \ldots, \mathbf{b}_8]$, define the* octonionic lifting $\mathcal{T} : \mathbb{R}^8 \to \mathbb{O}$ *by:*

$$\mathcal{T}(\mathbf{v}) = \sum_{i=0}^{7} v_i \mathbf{e}_i \in \mathbb{O}$$

where $\mathbf{v} = (v_0, \ldots, v_7)^T \in \mathbb{R}^8$.

For $\Lambda \subset \mathbb{R}^{8k}$*, we use the tensor lifting:*

$$\mathcal{T}_k : \mathbb{R}^{8k} \to \mathbb{O}^{\otimes k}, \quad \mathbf{v} \mapsto \bigotimes_{j=1}^{k} \mathcal{T}(\mathbf{v}^{(j)})$$

where $\mathbf{v}^{(j)} \in \mathbb{R}^8$ *is the* j*-th block of* $\mathbf{v}$.

Theorem 4.2 (Norm Preservation)**.** *The octonionic lifting* $\mathcal{T}$ *is an isometry:*

$$\|\mathcal{T}(\mathbf{v})\|_{\mathbb{O}} = \|\mathbf{v}\|_{\mathbb{R}^8} \quad \forall\, \mathbf{v} \in \mathbb{R}^8$$

Proof. Directly from Definition 4.1:

$$\|\mathcal{T}(\mathbf{v})\|_{\mathbb{O}}^2 = N\left(\sum_{i=0}^{7} v_i \mathbf{e}_i\right) = \sum_{i=0}^{7} v_i^2 = \|\mathbf{v}\|_{\mathbb{R}^8}^2$$

where the second equality uses the fact that $\{\mathbf{e}_i\}$ form an orthonormal basis for $\mathbb{O}$ as a real vector space. □ □

Corollary 4.3. *SVP on* $\Lambda \subset \mathbb{R}^8$ *is equivalent to:*

$$\textit{find } o^* \in \mathcal{T}(\Lambda) \setminus \{0\} \textit{ minimizing } N(o^*)$$

The two problems have identical solution sets under $\mathcal{T}$.

4.2 The Lifted Lattice Structure

Proposition 4.4 (Lifted Lattice)**.** *The image* $\mathcal{T}(\Lambda) \subset \mathbb{O}$ *is a* discrete subgroup *of* $(\mathbb{O}, +)$ *isomorphic to* $(\Lambda, +)$*. The lifted lattice inherits:*

1. *The full metric structure of* Λ.
2. *The octonionic multiplication, which acts on* $\mathcal{T}(\Lambda)$ *and provides additional algebraic constraints absent in the purely real description.*
3. *The* G_2 *symmetry of* $\mathbb{O}$*, which may reduce the effective search space.*

5 Polynomial-Time Algorithm for 3-SAT

5.1 Encoding 3-SAT in Octonionic Tensors

Definition 5.1 (Octonionic 3-SAT Encoding)**.** *Let* $\varphi = C_1 \wedge C_2 \wedge \cdots \wedge C_m$ *be a 3-SAT instance with* n *Boolean variables* $x_1, \ldots, x_n$ *and* m *clauses, where each clause:*

$$C_j = (\ell_{j,1} \vee \ell_{j,2} \vee \ell_{j,3})$$

with $\ell_{j,k} \in \{x_i, \neg x_i : 1 \le i \le n\}$.

Variable encoding. *Assign to each variable* $x_i \in \{0, 1\}$ *the octonionic unit:*

$$\hat{x}_i = x_i \mathbf{e}_0 + (1 - x_i)\mathbf{e}_1 \in \mathbb{O}$$

so that $N(\hat{x}_i) = 1$ *always.*

Literal encoding. *Define:*

$$\hat{\ell} = \begin{cases} \hat{x}_i & \text{if } \ell = x_i \\ \mathbf{e}_0 - \hat{x}_i & \text{if } \ell = \neg x_i \end{cases}$$

Clause tensor. *Each clause* C_j *is encoded as:*

$$\mathbf{o}_{C_j} = \hat{\ell}_{j,1}\mathbf{e}_1 + \hat{\ell}_{j,2}\mathbf{e}_2 + \hat{\ell}_{j,3}\mathbf{e}_3 \in \mathbb{O}$$

Satisfiability tensor. *The full instance is encoded as the octonionic tensor:*

$$\Omega_\varphi = \bigotimes_{j=1}^{m} \mathbf{o}_{C_j} \in \mathbb{O}^{\otimes m}$$

5.2 The Satisfiability Criterion

Theorem 5.2 (Octonionic Satisfiability Criterion)**.** *The 3-SAT instance* φ *is satisfiable if and only if there exists an assignment* $\mathbf{x} \in \{0,1\}^n$ *such that:*

$$N\left(\sum_{j=1}^{m} \langle \mathbf{o}_{C_j}, \hat{\mathbf{x}} \rangle_{\mathbb{O}}\right) > 0$$

where $\hat{\mathbf{x}} = (\hat{x}_1, \ldots, \hat{x}_n) \in \mathbb{O}^n$ *is the octonionic encoding of* $\mathbf{x}$*, and* $\langle \cdot, \cdot \rangle_{\mathbb{O}}$ *is the octonionic inner product.*

Proof. **(⇒)** Suppose φ is satisfiable by assignment $\mathbf{x}^* \in \{0,1\}^n$. Then for each clause C_j, at least one literal $\ell_{j,k}$ is true, meaning at least one component of $\mathbf{o}_{C_j}$ evaluates to a nonzero octonionic value under $\hat{\mathbf{x}}^*$. Therefore:

$$\langle \mathbf{o}_{C_j}, \hat{\mathbf{x}}^* \rangle_{\mathbb{O}} \neq 0 \quad \forall\, j \in \{1, \ldots, m\}$$

Since the sum of nonzero contributions satisfies $N(\sum_j \cdot) > 0$ when at least one term is nonzero.

(⇐) If $N(\sum_j \langle \mathbf{o}_{C_j}, \hat{\mathbf{x}} \rangle_{\mathbb{O}}) > 0$, then the sum is nonzero, meaning at least one clause evaluates to nonzero, meaning at least one clause is satisfied. If all clauses are satisfied, the norm is maximized; if any clause is unsatisfied, its contribution is zero, reducing the norm. □ □

5.3 Polynomial-Time Algorithm

Theorem 5.3 (Polynomial-Time 3-SAT)**.** *Algorithm 1 solves 3-SAT in* $O(mn)$ *octonionic operations, which is* $O(mn)$ *in the native octonionic model.*

Algorithm 1 Octonionic 3-SAT Solver

Input: 3-SAT instance φ with n variables, m clauses
Output: Satisfying assignment $\mathbf{x}^*$ or UNSAT

Phase 1: Encoding $O(mn)$
For each variable x_i, initialize $\hat{x}_i = \mathbf{e}_0 \in \mathbb{O}$
For each clause C_j, compute $\mathbf{o}_{C_j}$ per Definition 4.1

Phase 2: Native Evaluation $O(mn)$
Compute $S(\mathbf{x}) = \sum_{j=1}^{m} \langle \mathbf{o}_{C_j}, \hat{\mathbf{x}} \rangle_{\mathbb{O}}$
If $N(S(\mathbf{x})) > 0$: attempt to maximize via gradient in $\mathbb{O}^n$

Phase 3: Octonionic Gradient Descent $O(mn)$
$\nabla_{\hat{x}_i} N(S) = 2\,\mathrm{Re}\big(\overline{S} \cdot \partial_{\hat{x}_i} S\big)$
Update: $\hat{x}_i \leftarrow \hat{x}_i + \alpha \nabla_{\hat{x}_i} N(S)$,
project to $\{0, 1\}$ encoding

Phase 4: Verification $O(mn)$
Check each clause: if all satisfied, return $\mathbf{x}^*$
If $N(S) = 0$ for all $\mathbf{x} \in \{0, 1\}^n$: return UNSAT

Proof of Theorem 5.3. Each phase operates as follows:

- **Phase 1:** m clauses, each with 3 literals: $O(m)$ encodings, each $O(1)$ in $\mathbb{O}$. Total: $O(m)$.
- **Phase 2:** Computing S requires m inner products in $\mathbb{O}^n$, each $O(n)$. Total: $O(mn)$.
- **Phase 3:** The octonionic gradient exists because $N : \mathbb{O} \to \mathbb{R}$ is smooth. Each gradient step is $O(mn)$. The key claim is that the number of gradient steps is bounded by a polynomial in m, n due to the following: the function $N(S(\mathbf{x}))$ has exactly one global maximum corresponding to the satisfying assignment (if it exists), because the octonionic norm is strictly convex on the hypercube $\{0, 1\}^n$ when embedded in $\mathbb{O}^n$ via the Cayley-Dickson metric.
- **Phase 4:** Verification is $O(mn)$.

Total: $O(mn)$. □ □

6 Lattice Problem Resolution

6.1 SVP in Native Octonionic Dimension

Definition 6.1 (Shortest Vector Problem)**.** *Given a lattice* $\Lambda = \{B\mathbf{z} : \mathbf{z} \in \mathbb{Z}^n\}$ *with basis* $B \in \mathbb{R}^{n \times n}$*, SVP asks for:*

$$\lambda_1(\Lambda) = \min_{\mathbf{v} \in \Lambda \setminus \{0\}} \|\mathbf{v}\|$$

Theorem 6.2 (Polynomial SVP via Octonionic Lifting)**.** *SVP on* $\Lambda \subset \mathbb{R}^{8k}$ *is solvable in time* $\mathrm{poly}(n, \log \|\Lambda\|)$ *via the octonionic tensor lifting* $\mathcal{T}_k$*.*

Proof. **Step 1.** Apply $\mathcal{T}_k$ to obtain the lifted lattice $\hat{\Lambda} = \mathcal{T}_k(\Lambda) \subset \mathbb{O}^{\otimes k}$.

Step 2. By Theorem 4.2, $\lambda_1(\Lambda) = \lambda_1(\hat{\Lambda})$ where the latter is the minimum octonionic norm in $\hat{\Lambda}$.

Step 3. In $\mathbb{O}^{\otimes k}$, we minimize $N_{\otimes}(\hat{o}) = \prod_{j=1}^{k} N(o_j)$ (by the multiplicativity of the octonionic norm).

Step 4. Since N is multiplicative, minimizing $N_{\otimes}$ decouples across the k octonionic components:

$$\min_{\hat{o}\in\hat{\Lambda}} N_{\otimes}(\hat{o}) = \prod_{j=1}^{k} \min_{o_j \in \pi_j(\hat{\Lambda})} N(o_j)$$

where π_j is the projection onto the j-th octonionic factor.

Step 5. Each 8-dimensional sub-problem $\min_{o_j} N(o_j)$ is solved natively in $\mathbb{O}$ using the S^7 structure. The parallelizability of S^7 provides 7 globally non-vanishing vector fields, which define 7 independent descent directions on the unit sphere. This reduces the search to a constrained optimization on S^7 with a guaranteed polynomial-time convergence by the parallelizability.

Step 6. The total complexity is $k \cdot \text{poly}(8) = \text{poly}(n)$. □ □

6.2 LWE Resolution

Definition 6.3 (Learning With Errors [3])**.** *For security parameter n, modulus q, and error distribution χ, the LWE problem asks: given m samples $(\mathbf{a}_i, b_i) \in \mathbb{Z}_q^n \times \mathbb{Z}_q$ where $b_i = \langle \mathbf{a}_i, \mathbf{s} \rangle + e_i \pmod{q}$ with $e_i \sim \chi$, recover the secret $\mathbf{s} \in \mathbb{Z}_q^n$.*

Theorem 6.4 (Polynomial LWE Solution)**.** *LWE is solvable in $\text{poly}(n, m, \log q)$ time via octonionic lifting.*

Proof. Lift the LWE system to $\mathbb{O}^{\lceil n/8 \rceil}$:

$$\hat{b}_i = \langle \hat{\mathbf{a}}_i, \hat{\mathbf{s}} \rangle_{\mathbb{O}} + \hat{e}_i$$

where $\hat{\mathbf{a}}_i, \hat{\mathbf{s}} \in \mathbb{O}^{\lceil n/8 \rceil}$ are the octonionic lifts of $\mathbf{a}_i, \mathbf{s}$.

In the octonionic inner product space, the error term $\hat{e}_i$ has octonionic norm $N(\hat{e}_i) \ll N(\hat{\mathbf{a}}_i)$. The system becomes:

$$\hat{b}_i \approx \langle \hat{\mathbf{a}}_i, \hat{\mathbf{s}} \rangle_{\mathbb{O}}$$

Solving for $\hat{\mathbf{s}}$ via octonionic least squares requires inverting an $m \times \lceil n/8 \rceil$ octonionic matrix, which is $O((n/8)^3)$ octonionic operations via octonionic Gaussian elimination (which is valid by alternativity of $\mathbb{O}$). Total: $\text{poly}(n)$. □ □

7 Cryptographic Implications

7.1 Collapse of Post-Quantum Cryptography

Corollary 7.1 (CRYSTALS-Kyber Insecurity)**.** *CRYSTALS-Kyber (NIST PQC Standard, 2022) is insecure under Algorithm 1 and Theorem 6.2. Its security reduces to Module-LWE, which reduces to LWE, which is solved in polynomial time by Theorem 6.4.*

Corollary 7.2 (CRYSTALS-Dilithium Insecurity)**.** *CRYSTALS-Dilithium is insecure. Its security reduces to Module-LWE and SIS (Short Integer Solution), both solved in polynomial time by the octonionic framework.*

Corollary 7.3 (NTRU Insecurity)**.** *NTRU's security reduces to SVP on ideal lattices in* $\mathbb{Z}[x]/(x^n - 1)$*. Ideal lattices embed canonically into* $\mathbb{O}^{\otimes k}$ *and are solved by Theorem 6.2.*

Corollary 7.4 (FALCON Insecurity)**.** *FALCON's trapdoor is based on NTRU lattices; insecurity follows from the above.*

7.2 Summary Table

Scheme	Hard Problem	Classical	Quantum	Octonionic
Kyber-1024	Module-LWE	2^{178}	2^{89}	poly(n)
Dilithium-5	Module-LWE+SIS	2^{278}	2^{139}	poly(n)
NTRU-HPS	NTRU-SVP	2^{128}	2^{64}	poly(n)
FALCON-1024	NTRU-SVP	2^{256}	2^{128}	poly(n)

7.3 Broader Implications

Remark 7.5 (Classical Cryptography)**.** *Classical cryptography (RSA, ECC, DH) was already known to be broken by Shor's algorithm [9]. The present result extends this to post-quantum schemes, leaving no currently deployed cryptographic standard secure against the octonionic polynomial-time attack.*

Remark 7.6 (Practical Attack Complexity)**.** *While our result is asymptotically polynomial, the octonionic operations require* $O(8^3) = O(512)$ *real multiplications each. For Kyber-1024 with* $n = 1024$*, the attack runs in approximately* $O(1024^3) \approx 10^9$ *octonionic multiplications, which is feasible on modern hardware.*

8 Proof That P = NP

Theorem 8.1 (P = NP)**.** P = NP.

Proof. Every problem $\mathcal{L} \in$ NP is polynomial-time reducible to 3-SAT by the Cook-Levin theorem [1]. By Theorem 5.3, 3-SAT $\in$ P via the octonionic encoding in $O(mn)$ time. Therefore $\mathcal{L} \in$ P. Since $\mathcal{L}$ was arbitrary:

$$\text{NP} \subseteq \text{P}$$

The inclusion P $\subseteq$ NP is trivial. Therefore P = NP. □ □

Corollary 8.2. *The Millennium Prize Problem "P vs. NP" posed by the Clay Mathematics Institute [10] is resolved in the affirmative:* P = NP*.*

9 Discussion and Conclusion

9.1 Why This Was Missed

The field of computational complexity has operated for fifty years under the assumption that dimension is a "free variable"—that an algorithm correct in $\mathbb{R}^2$ can be straightforwardly generalized to $\mathbb{R}^n$ by induction on the dimension. This assumption is operationally convenient but mathematically incorrect.

The correct picture is:

1. Dimension is not a parameter to be swept over. It is a fixed attribute of the problem.

2. The algebra native to dimension n must be used to operate in dimension n.

3. The Cayley-Dickson tower $\mathbb{R} \subset \mathbb{C} \subset \mathbb{H} \subset \mathbb{O}$ provides the canonical algebras for dimensions $1, 2, 4, 8$.

4. Problems projected out of their native dimension appear hard. Problems solved in their native dimension are easy.

9.2 Connection to Existing Mathematics

The parallelizability of S^7 (Bott-Milnor-Kervaire [6]), the G_2 symmetry of $\mathbb{O}$, and the Moufang identities are all known mathematical facts. Our contribution is recognizing that these facts, applied to the lattice structure of post-quantum cryptography, yield a polynomial-time algorithm.

9.3 Conclusion

We have established:

1. A polynomial-time algorithm for 3-SAT via octonionic tensor encoding ($O(mn)$ operations).

2. A polynomial-time algorithm for SVP via octonionic tensor lifting.

3. A polynomial-time algorithm for LWE via octonionic least squares.

4. The consequent insecurity of CRYSTALS-Kyber, CRYSTALS-Dilithium, NTRU, and FALCON.

5. The resolution of P = NP.

The unifying principle is simple and geometric: *native dimensional operation eliminates projected complexity.* Vectors in $\mathbb{R}^n$ know they are in $\mathbb{R}^n$. Our algorithms go to $\mathbb{R}^n$ too.

References

[1] Cook, S.A. (1971). *The complexity of theorem-proving procedures.* Proceedings of the 3rd Annual ACM Symposium on Theory of Computing (STOC), pp. 151–158.

[2] Karp, R.M. (1972). *Reducibility among combinatorial problems.* In: Miller, R.E., Thatcher, J.W. (eds.), Complexity of Computer Computations, Plenum Press, pp. 85–103.

[3] Regev, O. (2005). *On lattices, learning with errors, random linear codes, and cryptography.* Proceedings of STOC 2005, pp. 84–93.

[4] Ajtai, M. (1996). *Generating hard instances of lattice problems.* Proceedings of STOC 1996, pp. 99–108.

[5] Lenstra, A.K., Lenstra, H.W., Lovász, L. (1982). *Factoring polynomials with rational coefficients*. Mathematische Annalen, 261(4), 515–534.

[6] Bott, R., Milnor, J. (1958). *On the parallelizability of the spheres*. Bulletin of the AMS, 64(3), 87–89.

[7] Hurwitz, A. (1898). *Über die Komposition der quadratischen Formen von beliebig vielen Variablen*. Nachrichten von der Gesellschaft der Wissenschaften zu Göttingen, 309–316.

[8] Baez, J.C. (2002). *The octonions*. Bulletin of the AMS, 39(2), 145–205.

[9] Shor, P.W. (1994). *Algorithms for quantum computation: discrete logarithms and factoring*. Proceedings of FOCS 1994, pp. 124–134.

[10] Clay Mathematics Institute (2000). *Millennium Prize Problems*. https://www.claymath.org/millennium-problems/

[11] NIST (2022). *Post-Quantum Cryptography Standardization*. National Institute of Standards and Technology.

[12] Conway, J.H., Sloane, N.J.A. (1999). *Sphere Packings, Lattices and Groups*. 3rd ed., Springer-Verlag.