

# The Kolmogorov Complexity of Mathematical Proof: A Self-Measuring Tower

Richard Hoekstra

cs.LO / math.LO – April 2026

---

## Abstract

We measure the Kolmogorov complexity of machine-verified proofs by type-erasing Lean 4 proof terms into TLC (a seven-constructor Kraft-saturating ternary lambda calculus) and applying DAG-CSE via hash-consing. The result is an exact, reproducible complexity measure over proof objects.

$K(\text{epsilon}^2 = 0) = 5$  DAG nodes exactly (upper = lower bound, proved in Lean).  $K(\text{dual\_mul}) = 1018$  nodes (292x compressed from 298,003 raw). The axiom is incompressible: it shares zero mutual information with every other theorem. Ring normalization is 99.66% redundant.

K grows sublinearly with proof depth ( $\alpha < 1$ ), with an information production rate of approximately 29 nodes per proof step. The optimal axiom by MDL is not the simplest ( $K = 5$ ) but the richest ( $K = 1018$ ): conditioning on `dual_mul` leaves only 217 nodes for the rest of the tower, versus 1230 when conditioning on `eps_sq`.

The self-measurement tower stabilizes at Level 2-3, with total reflection cost 136 nodes = 27x the axiom. Meta-measurement is smaller than measurement:  $R(1 \rightarrow 2) = 0.3x$ .  $K = 5$  is universal for all  $x^2 = c$  equations (nilpotent, involution, idempotent), with a strict gap to  $K \geq 7$  for anything more complex. Order 2 is the K-optimal algebraic structure.

Among confluence proofs for order-2 algebras,  $K(\text{CR\_idempotent}) = 273 < K(\text{CR\_nilpotent}) = 444 < K(\text{CR\_involution}) = 578$ . Nilpotent is not the simplest confluence proof – idempotent is.

Appendix A presents a complementary approach: a formalized PPM-C compressor in Lean 4 with exact rational arithmetic, whose roundtrip property is proved with zero sorry's, yielding the first machine-verified Kolmogorov complexity upper bound on a nontrivial string.

---

## 1. Introduction: Measuring Proof Complexity in Bits

How complex is a mathematical proof? Not in the sense of difficulty (which is psychological) or length (which depends on notation), but in the information-theoretic sense: how many bits are required to specify the proof uniquely?

Kolmogorov complexity  $K(x)$  is the length of the shortest program that outputs  $x$ . For natural-language proofs,  $K$  is uncomputable. But machine-verified proofs are concrete objects – Lean 4 produces explicit proof terms that can be inspected, serialized, and measured. The question becomes tractable: given a proof term, what is its Kolmogorov complexity relative to a fixed universal encoding?

We propose and implement a complete pipeline:

1. **Type erasure.** `Lean.Expr` (the proof term) is converted to TLC, a seven-constructor ternary lambda calculus, by forgetting all type information. The resulting term captures the computational content – the “what” of the proof – without the “why.”
2. **DAG-CSE.** The TLC tree is hash-consed into a directed acyclic graph. Structurally identical subtrees are shared. The DAG node count is our complexity measure:  $K\_DAG(\text{proof}) =$  number of unique subtrees.
3. **Exact bounds.** For the simplest proofs, we prove upper = lower bounds on  $K\_DAG$ , yielding exact Kolmogorov complexities. For larger proofs, the DAG provides tight upper bounds with compression ratios up to 292x.

We apply this to the epsilon tower: nine theorems about nilsquare elements ( $\epsilon^2 = 0$ ) in commutative rings, ranging from the 5-node axiom to the 298,003-node ring normalization proof. The numbers reveal structure invisible to informal mathematics.

---

## 2. The TLC Bridge

### 2.1 Seven Constructors

TLC (Ternary Lambda Calculus) has seven constructors encoded over a ternary alphabet:

Code	Constructor	Trit cost	Pre-CSE freq	Role
0	<code>app(fn, arg)</code>	1 trit	49%	Application
10	<code>const(idx)</code>	2+ trits	27%	Global reference
11	<code>var(idx)</code>	2+ trits	25%	Local (de Bruijn)
12	<code>lam(body)</code>	2+ trits	2%	Abstraction
20	<code>let(val, body)</code>	2+ trits	0%	Sharing
21	<code>case(scrut, alts)</code>	2+ trits	0%	Observation
22	<code>quote(term)</code>	2+ trits	0%	Reification

The Kraft sum is:

$$1/3 + 6/9 = 1/3 + 2/3 = 1$$

Saturating. No prefix code over ternary can use these seven slots more efficiently. The encoding is tight.

The first trit partitions the universe: 0 = application (data), 1 = binding/reference (code), 2 = meta. The second trit within each partition distinguishes global/local/binder (for 1x) and share/observe/reify (for 2x).

### 2.2 Type Erasure: `Lean.Expr` to TLC

The bridge function `exprToTLC` maps Lean’s kernel expressions to TLC by forgetting types:

```
Lean.Expr.app f a      -> TLC.app (exprToTLC f) (exprToTLC a)
Lean.Expr.const name _ -> TLC.const 0
Lean.Expr.bvar i       -> TLC.var i
Lean.Expr.lam _ _ b _  -> TLC.lam (exprToTLC b)
```

```

Lean.Expr.forallE _ _ b _ -> TLC.lam (exprToTLC b)
Lean.Expr.letE _ _ v b _ -> TLC.letb (exprToTLC v) (exprToTLC b)
Lean.Expr.mdata _ e      -> exprToTLC e    (transparent)
Lean.Expr.proj _ _ e     -> TLC.case (exprToTLC e) (TLC.var 0)
Lean.Expr.sort _        -> TLC.quote (TLC.var 0)
_                       -> TLC.var 0      (fallback)

```

Type information (in `lam`, `forallE`, `letE`) is discarded. The resulting TLC term is the computational skeleton of the proof. This erasure is partial but honest: the `exprToTLC` function is itself `partial` in Lean (it recurses without a termination proof), and when we measure the instrument later, this partiality shows up as opacity.

## 2.3 DAG-CSE via Hash-Consing

The key compression step is DAG-CSE (Common Subexpression Elimination). The TLC tree is traversed depth-first; each subtree is hashed (via Lean’s built-in `Hashable` instance) and checked against a `Std.HashSet`. If already seen, the subtree is shared rather than duplicated.

The result is a DAG structure (defined in `TLCCompress.lean`): - `nodes` : `Array Node` – unique subexpressions in topological order - `root` : `NodeId` – the root of the expression - `uses` : `Array Nat` – reference count for each node

A `Node` is a TLC constructor with children replaced by `NodeId` indices into the array. The `DAG.nodeCount` (= `nodes.size`) is our `K_DAG` measure.

This is exact: no approximation, no sampling. The DAG is deterministic for a given TLC term. Two terms have the same `K_DAG` if and only if they have the same set of unique subtrees.

## 3. K of the Epsilon Tower

### 3.1 Per-Theorem Measurements

We measure eight theorems from the epsilon tower – proofs about nilsquare elements ( $\epsilon^2 = 0$ ) in commutative rings:

Theorem	Raw nodes	DAG nodes	Shared	Ratio	Layer
<b>eps_sq</b>	6	<b>5</b>	1	<b>1x (exact)</b>	L1
eps_nilpotent	111	38	13	2x	L1
eps_zero_divisor	307	57	26	5x	L1
mulEps_mulEps	567	53	38	10x	L2
mulEps_comp_zero	1,964	80	60	24x	L2
eps_pow_two	1,452	100	55	14x	L1
exact_iff	553	115	56	4x	L2
<b>dual_mul</b>	<b>298,003</b>	<b>1,018</b>	<b>527</b>	<b>292x</b>	L1

The range is dramatic: from 5 to 1,018 DAG nodes, spanning three orders of magnitude in compression ratio.

### 3.2 $K(\text{eps\_sq}) = 5$ Exactly

This is the first exact Kolmogorov complexity of a mathematical proof.

**Upper bound.** We construct the DAG explicitly. The proof term for  $\text{epsilon}^2 = 0$  is:

```
lam(app(app(const 22, var 0), var 0))
```

meaning “lambda epsilon. mul(epsilon, epsilon).” The five DAG nodes are:

```
Node 0: const 22      (the multiplication operator)
Node 1: var 0         (epsilon -- shared, referenced twice)
Node 2: app(0, 1)    (mul applied to epsilon)
Node 3: app(2, 1)    (mul(epsilon) applied to epsilon -- reuses Node 1)
Node 4: lam(3)       (the proof binder)
```

DAG size = 5 (proved: `epsSqProof_dag_nodes` in `KLowerBound.lean`).

**Lower bound.** All five nodes are structurally distinct. Specifically: - Node 0 is `const 22` (a leaf) - Node 1 is `var 0` (a different leaf) - Node 2 is `app(0, 1)` (references two distinct children) - Node 3 is `app(2, 1)` (references different first child than Node 2) - Node 4 is `lam(3)` (a unary node, distinct from all binary/leaf nodes)

No further sharing is possible. The lower bound is 5 (proved: `epsSqProof_dag_nodes_distinct` in `KLowerBound.lean`).

**Combined.**  $K_{\text{DAG}}(\text{eps\_sq}) = 5$  exactly (`K_epsSq_tight`). The axiom is its own shortest description.

### 3.3 $K(\text{dual\_mul}) = 1,018$ : Ring Normalization is 99.66% Redundant

The dual multiplication theorem (`dual_mul`: the associativity-distributivity property) has 298,003 raw nodes but only 1,018 unique subtrees. Lean’s `ring` tactic generates massive proof terms by unfolding ring axioms mechanically. The proof is syntactically enormous but structurally repetitive.

Redundancy =  $1 - (1,018 / 298,003) = 99.66\%$ .

Of the 1,018 unique nodes, 527 (51.8%) are shared – referenced more than once in the original tree. The compression ratio of 292x is the highest in the tower.

### 3.4 Pre-CSE Constructor Frequencies

Across the tower, the pre-CSE (tree) constructor distribution is:

Constructor	Pre-CSE frequency
<code>app</code>	49%
<code>const</code>	27%
<code>var</code>	25%
<code>lam</code>	2%
<code>let, case, quote</code>	~0%

After DAG-CSE, the distribution shifts dramatically:

Constructor	Post-CSE frequency
app	94%
var	3%
lam	3%
const	~0%
let, case, quote	0%

The leaves (const, var) collapse into shared references; the structural `app` nodes that wire them together remain unique. Shannon entropy drops from 0.948 to 0.283 trits. Mathematics, after deduplication, is 94% application – applying things to other things.

## 4. Tower Cohesion

### 4.1 Combined DAG

All nine Layer 1 + Layer 2 theorems are merged into a single hash-consed DAG:

Metric	Value
Sum of individual DAGs	1,571 nodes
Combined DAG	1,409 nodes
Cross-theorem shared nodes	162 nodes
<b>Cohesion</b>	<b>10%</b>

The tower is 90% modular (each theorem carries its own machinery) and 10% cohesive (shared mathematical vocabulary). `dual_mul` dominates at 1,018 of 1,571 individual DAG nodes (65%).

### 4.2 Mutual Information Matrix

Mutual information  $I(A; B) = K(A) + K(B) - K(A, B)$ , computed via combined DAG:

**Most related pairs:**

Rank	Pair	I
1	<code>eps_pow_two &lt;-&gt; eps_zero_divisor</code>	31
2	<code>eps_nilpotent &lt;-&gt; eps_pow_two</code>	22
3	<code>mulEps_mulEps &lt;-&gt; exact_iff</code>	19

**Most independent:**

Pair	I
<code>eps_sq &lt;-&gt; everything</code>	<b>0</b>

eps\_sq has ZERO mutual information with every other theorem. It is the pure atomic axiom – it shares no proof structure with anything. It tells you something you cannot learn from any other theorem’s proof term.

**Hub theorem:** eps\_pow\_two, with total I = 104 (sum of pairwise MI with all other theorems). Average I per pair = 11.3.

The hub is not the axiom (which is isolated) nor the largest theorem (which is too specialized). It is the intermediate theorem that connects the most proof machinery.

## 5. K vs. Proof Depth

### 5.1 Sublinear Growth

Assigning logical dependency depth to each theorem (depth 0 = axiom, depth n = depends on theorems of depth < n):

Depth	Avg DAG K	Representative theorems
0	5	eps_sq (the axiom)
1	81	eps_nilpotent, eps_zero_divisor, eps_pow_two
2	370	dual_mul (outlier: ring tactic), mulEps
3	82	mulEps compositions, exact_iff
4	123	derivation theorems
5	271	deformation theorems

Linear regression across all 17 measured theorems gives:

$dK/d(\text{depth}) \sim 29$  DAG nodes per proof step

The power-law exponent  $\alpha < 1$ : K grows sublinearly with depth. The information production rate DECREASES as proofs go deeper. This is the signature of increasing structural reuse. Deeper proofs build on more shared vocabulary, so each new step adds fewer truly new subtrees.

### 5.2 Interpretation

The sublinearity has a direct mathematical meaning: composition is cheaper than construction. The master theorem that assembles lemmas (NilSquare.lean, density 3.3 nodes/line) is 6x more compact than the lemmas themselves (Epsilon.lean, density 19.3 nodes/line). Assembly is cheap. Proving is expensive.

## 6. Optimal Axiom Selection

### 6.1 The MDL Criterion

Which single theorem, if taken as axiom, minimizes the total Kolmogorov complexity of describing the tower? This is the Minimum Description Length (MDL) criterion applied to proof selection.

For each candidate axiom A: -  $K(A)$  = DAG nodes of A alone -  $K(\text{rest} \mid A)$  = combined DAG of A + rest, minus  $K(A)$  - Total =  $K(A) + K(\text{rest} \mid A) = K(\text{combined})$

The combined DAG is constant across candidates, so the ranking reduces to: which axiom contributes the most DAG nodes to the combined structure?

## 6.2 Results

Candidate	$K(\text{axiom})$	$K(\text{rest} \mid A)$	Winner?
<b>dual_mul</b>	1,018	<b>217</b>	<b>YES</b>
exact_iff	115	1,120	
eps_pow_two	100	1,135	
eps_sq	5	1,230	

**dual\_mul wins.** The AD property is the most information-efficient starting point – not because it is simple (it is the largest theorem) but because its proof term contains 83% of the tower’s shared algebraic machinery. The simplest axiom (eps\_sq,  $K = 5$ ) leaves the highest conditional cost ( $K = 1,230$ ). The richest axiom (dual\_mul,  $K = 1,018$ ) leaves the lowest ( $K = 217$ ).

## 6.3 Fundamentality != Simplicity

This is the central insight of the section. In traditional mathematics, the “fundamental” theorem is usually the simplest – the one from which everything else follows with minimal effort. But  $K$ -complexity reveals a different structure: the most fundamental theorem (in the MDL sense) is the one whose proof term contains the most reusable structure.

dual\_mul’s ring normalization proof is massive and repetitive, but that very repetitiveness means it encodes the ring axioms in every possible combination. Conditioning on it, the rest of the tower becomes nearly trivial.

eps\_sq is the opposite: incompressible, isolated ( $I = 0$  with all), and informationally atomic. It is the purest axiom but the worst foundation for reconstruction.

## 7. Self-Measurement

### 7.1 The Instrument Measures Itself

The measurement pipeline (TLCBridge + TLCCompress) is itself a collection of Lean definitions. We measure them with their own tools:

Function	$K$ (trits)	DAG nodes	Role
exprToTLC	16	4	Type erasure (partial – opaque)
tritLength	749	44	Trit counting
nodeCount	659	42	Node counting
getConstValue	101	14	Extract proof term
<b>measure</b>	<b>1,132</b>	<b>83</b>	The full instrument
dagCSE	232	30	Hash-consed compression

$K(\text{measure}) = 83$  DAG nodes. The measurement overhead is  $K(\text{measure}) / K(\text{eps\_sq}) = 83 / 5 = 16.6x$  (or  $1,132 / 16 = 70x$  in raw trits).

The instrument is 16.6x more complex (in DAG nodes) than the simplest thing it measures. Note the opacity: `exprToTLC` is **partial** in Lean, so its recursive kernel is wrapped in an opaque thunk. The snake can measure its body (1,132 trits) but not its eyes.

## 7.2 The Fixed-Point Tower

The self-reference chain computes  $K$  at each level of measurement:

```
Level 0: K(eps_sq)      = 5 nodes (the axiom)
Level 1: K(measure)    = 83 nodes (measures Level 0)
Level 2: K(meta-measure) = 23 nodes (measures Level 1)
Level 3: K(meta-meta)  = 25 nodes (measures Level 2)
```

Reflection overhead  $R(n \rightarrow n+1) = K(n+1) / K(n)$ :

```
R(0->1) = 83/5   = 16.6x (first instrument is 16.6x the axiom)
R(1->2) = 23/83  = 0.3x (meta-measurement is SMALLER)
R(2->3) = 25/23  = 1.1x (near-isomorphic -- stabilized)
```

## 7.3 The Tower Stabilizes

The tower stabilizes at Level 2-3. Measuring a measurement has the same complexity as measuring a measurement of a measurement. The  $R(2 \rightarrow 3) = 1.1x$  ratio is near unity; the levels become structurally isomorphic.

The critical observation is  $R(1 \rightarrow 2) < 1$ : meta-measurement is SMALLER than measurement. This is because the meta-level functions (`SelfMeasure.rawTrits`, `SelfMeasure.dagNodes`) are thin wrappers that call `TLCBridge.measure`. They strip presentation overhead. Going up the tower simplifies the object being measured.

## 7.4 The Price of Reflection

The Chaitin constant – the total information cost of the measurement tower:

```
5 + 83 + 23 + 25 = **136 DAG nodes**
```

The total price of self-awareness is 136 nodes = 27x the axiom. This is finite, bounded, and exactly computable. The system knows exactly how much it costs to know itself.

---

## 8. $K = 5$ Universality

### 8.1 All $x^2 = c$ Axioms Have $K = 5$

The three fundamental order-2 axiom systems: - **Nilpotent:**  $\epsilon * \epsilon = 0$  (NilSquare rings)  
- **Involution:**  $j * j = 1$  (involutive elements) - **Idempotent:**  $e * e = e$  (band semigroups)

all share the identical TLC proof skeleton:

```
lam(app(app(const k, var 0), var 0))
```

differing only in the constant index  $k$  (encoding the operation) and the right-hand side constant. The DAG structure is isomorphic across all three:

```

Node 0: const k      (the operation)
Node 1: var 0        (the element -- shared)
Node 2: app(0, 1)    (partial application)
Node 3: app(2, 1)    (full application, reuses Node 1)
Node 4: lam(3)       (the proof binder)

```

$K\_DAG = 5$  for all three. Proved in Lean for all constant/variable index pairs in  $\text{Fin}(10) \times \text{Fin}(10)$  (`K_DAG_5_forall_small_indices` in `KFiveUniversal.lean`).

## 8.2 The Gap: $K \geq 7$ for Higher-Order

More complex equations require strictly more DAG nodes:

Equation	TLC encoding	Tree nodes	DAG nodes
$x^2 = c$	<code>lam(app(app(mul, x), x))</code>	6	<b>5</b>
$x^3 = c$	<code>lam(app(app(mul, x^2), x))</code>	8	<b>7</b>
$(x^2)^2 = c$	<code>lam(app(app(mul, x^2), x^2))</code>	10	<b>7</b>
$x^2 + x = c$	<code>lam(app(app(add, x^2), x))</code>	8	<b>8</b>
$x^2 + x^2 = c$	<code>lam(app(app(add, x^2), x^2))</code>	10	<b>8</b>

There is a strict gap of 2 between  $x^2 = c$  ( $K = 5$ ) and any more complex equation ( $K \geq 7$ ). No equation with  $K = 6$  exists in this family. Proved: `K_DAG_gap` in `KFiveUniversal.lean`.

## 8.3 Order 2 is K-Optimal

Why the gap? The  $x^2 = c$  form is the unique equation that: 1. Uses exactly one binary operation (one `const` node) 2. Has variable reuse (one `var` node, shared between both operand positions) 3. Requires exactly two application nodes to wire the operation to its arguments 4. Needs exactly one lambda for the proof binder

Any additional operation (e.g., `add` in  $x^2 + x$ ) introduces a new `const` node and new `app` nodes. Any higher power ( $x^3, x^4$ ) introduces additional `app` layers. The minimum is achieved if and only if the equation has the form  $\text{op}(x, x) = c$  – a single binary operation with shared argument. This is precisely the order-2 condition.

**Order 2 is the K-optimal algebraic structure.** This is not a convention but a theorem about information content.

# 9. Confluence Comparison

## 9.1 Three Order-2 Confluence Proofs

We measure the Kolmogorov complexity of confluence proofs for the three order-2 algebras. The confluence question: if two elements of the given type combine, does the result commute?

- **Nilpotent confluence** (`nilsquare_comm_zero`): if  $\epsilon_1^2 = 0$  and  $\epsilon_2^2 = 0$  and  $(\epsilon_1 + \epsilon_2)^2 = 0$ , then  $\epsilon_1 * \epsilon_2 = 0$ .
- **Involution confluence** (`CR_involution`): if  $j_1^2 = 1$  and  $j_2^2 = 1$  and  $(j_1 * j_2)^2 = 1$ , then  $j_1 * j_2 = j_2 * j_1$ .
- **Idempotent confluence** (`CR_idempotent`): if  $e_1^2 = e_1$  and  $e_2^2 = e_2$ , then  $e_1 * e_2 * e_1 = e_1 * e_2$ .

## 9.2 K\_DAG Measurements

Confluence proof	K_DAG (nodes)
<b>CR_idempotent</b>	<b>273</b>
CR_nilpotent	444
CR_involution	578

**Ordering: K(idempotent) < K(nilpotent) < K(involution).**

## 9.3 Nilpotent is NOT the Simplest Confluence

The prediction was that nilpotent (the axiom with  $K = 5$  for the base equation) would also have the simplest confluence proof. This is falsified.

The idempotent confluence proof is simplest ( $K = 273$ ) because the band equation  $e_1 * e_2 * e_1 = e_1 * e_2$  follows from commutativity plus one application of  $e_1^2 = e_1$ , using ring to rearrange. The nilpotent proof requires the hypothesis that  $(\epsilon_1 + \epsilon_2)^2 = 0$  and an invertibility condition on 2, making the proof structurally more complex. The involution proof requires group-theoretic reasoning (inverses), which introduces the most structural overhead.

This demonstrates that K-complexity of an axiom (all three have  $K = 5$ ) does not predict K-complexity of its derived theorems. Simplicity of the axiom and simplicity of the theory are different measures.

# 10. Discussion: What the Numbers Say

## 10.1 Seven Numbers

The complete K-portrait of the  $\epsilon^2 = 0$  theory can be summarized in seven numbers:

<code>K(eps_sq)</code>	=	5	(the axiom -- incompressible, I=0 with all)
<code>K(exact_iff)</code>	=	115	(the heart -- 4x compressed, hub-adjacent)
<code>K(dual_mul)</code>	=	1,018	(the richest -- optimal axiom, 292x compressed)
<code>K(measure)</code>	=	83	(the instrument -- 16.6x the axiom)
<code>K(meta)</code>	=	23	(the meta -- 0.3x the instrument)
<code>K(tower)</code>	=	1,409	(all together -- 10% cohesion)
<code>K(reflection)</code>	=	136	(the price of self-awareness -- 27x the axiom)

## 10.2 Axioms are Incompressible

$K(\text{eps\_sq}) = 5$  exactly, with upper = lower bound proved in Lean. No shorter description of this proof exists. The axiom IS the information. Moreover, it has zero mutual information with every other theorem – it is informationally atomic, sharing no proof structure with anything else in the tower.

## 10.3 Ring Normalization is 99.66% Redundant

`dual_mul` has 298,003 raw nodes but only 1,018 unique subtrees. Lean’s `ring` tactic generates massive but repetitive proof terms. The tactic’s strategy is brute-force normalization: expand everything, then compare normal forms. This is sound but profligate. The DAG reveals that 99.66% of the proof term is structural repetition.

## 10.4 The Theory is 10% Cohesive

Cross-theorem sharing saves 162 nodes out of 1,571 individual DAG nodes. The theorems share vocabulary (ring operations, epsilon, basic lemmas) but not machinery (each proof carries its own normalization and rewriting infrastructure). Mathematics, at the proof-term level, is overwhelmingly modular.

## 10.5 Assembly is 6x More Compact Than Proof

The information density measurements reveal a clear pattern:

File	DAG nodes	Lines	Density
<code>NilSquare.lean</code>	295	90	3.3 n/l
<code>EpsilonModule.lean</code>	450	99	4.5 n/l
<code>Nilpotent.lean</code>	411	32	12.8 n/l
<code>EpsilonVariations.lean</code>	2,162	113	19.1 n/l
<code>Epsilon.lean</code>	1,811	94	19.3 n/l

The master theorem file (`NilSquare.lean`) that assembles lemmas is 6x more compact than the files containing the lemmas themselves. Composing is cheaper than constructing.

## 10.6 Application Dominates After CSE

Post-CSE, application nodes constitute 94% of all unique DAG nodes. Abstraction (`lam`) is 3%. Sharing (`let`) is 0%. Mathematics, at its structural core, is applying things to other things. The leaves (constants and variables) are few and shared; the wiring that connects them is unique and dominant.

## 10.7 The Measurement Paradox

The meta-measurement (Level 2) is SMALLER than the measurement (Level 1):  $R(1 \rightarrow 2) = 23/83 = 0.3x$ . This seems paradoxical – shouldn’t measuring a measurement add complexity? No. The meta-level functions are thin wrappers. They call the same measurement infrastructure but on a simpler object (a measurement function rather than a proof term). Going up the tower strips overhead rather than adding it.

This resolves a philosophical puzzle about self-reference: the infinite regress of “measuring the measurer” does not diverge. It converges, and it converges fast (by Level 2–3).

## 10.8 Order 2 is Special

The  $K = 5$  universality result and the  $K \geq 7$  gap theorem together identify order 2 as the information-theoretically optimal algebraic structure. This is not a consequence of any particular axiom system but of the structure of proof terms themselves. Any equation of the form  $x * x = c$  requires exactly 5 DAG nodes. Any equation involving higher powers, additional operations, or non-shared variables requires strictly more.

The gap of 2 (from  $K = 5$  to  $K \geq 7$ ) is a structural discontinuity: there exists no algebraic equation with  $K = 6$  in the family of polynomial-identity equations.

---

## 11. Conclusion

We have measured the Kolmogorov complexity of machine-verified mathematical proofs using a concrete, reproducible pipeline: Lean proof terms  $\rightarrow$  TLC type erasure  $\rightarrow$  DAG-CSE hash-encoding  $\rightarrow$  node count. The method yields exact values for small proofs and tight upper bounds (with compression ratios up to 292x) for large ones.

The results expose structure invisible to informal mathematical reasoning:

1. **The axiom  $\epsilon^2 = 0$  has  $K = 5$  exactly** – an exact Kolmogorov complexity, proved with matching upper and lower bounds. The axiom is incompressible and informationally atomic (zero mutual information with all other theorems).
2. **The richest theorem is the best axiom.** `dual_mul` ( $K = 1,018$ ) minimizes the conditional complexity of the rest of the tower. Fundamentality is not simplicity but structural centrality.
3. **The self-measurement tower stabilizes at Level 2–3** with total cost 136 nodes. Self-awareness has a finite, measurable price. Meta-measurement is smaller than measurement.
4.  **$K = 5$  is universal for order-2 equations** (nilpotent, involution, idempotent). A strict gap of 2 separates order-2 from everything more complex. Order 2 is the  $K$ -optimal algebraic structure.
5. **Confluence complexity does not follow axiom complexity.** Idempotent confluence ( $K = 273$ ) is simpler than nilpotent ( $K = 444$ ), despite the axioms having identical  $K = 5$ .
6. **Ring normalization is 99.66% redundant.** After DAG-CSE, Lean’s `ring` tactic proofs compress by 292x.
7. **Mathematics is 94% application.** After deduplication, the unique structure of proof terms is overwhelmingly the wiring (app nodes) that connects shared leaves.

The snake measures its own length. The length is 83 nodes. The price of knowing this is 136.

## References

### Lean Source Files

```
Proof/TLC.lean          -- TLC term definition (7 constructors, Kraft-saturating)
Proof/TLCBridge.lean   -- Lean.Expr -> TLC type erasure, K-measurement
Proof/TLCCompress.lean -- DAG-CSE via hash-consing
Proof/KLowerBound.lean -- K(eps_sq) = 5 exactly (upper = lower)
Proof/KFiveUniversal.lean -- K=5 universal for x^2=c, gap to K>=7
Proof/KStructure.lean  -- K vs depth (sublinear, alpha < 1)
Proof/OptimalAxiom.lean -- dual_mul wins MDL selection
Proof/MutualInfo.lean  -- I(A;B) pairwise matrix, eps_sq atomic
Proof/TowerMeasure.lean -- Cross-theorem cohesion (10%)
Proof/PostCSEFrequency.lean -- Post-CSE distribution (app = 94%)
Proof/SelfMeasure.lean -- K(measure) = 83
Proof/FixedPoint.lean  -- Self-reference tower stabilizes at 23-25
Proof/ConfluenceComparison.lean -- K(CR) for three order-2 algebras
Proof/KOLMOGOROV.md    -- Measurement data and analysis

Proof/PPMBayes.lean    -- Bayesian prediction framework (1 sorry)
Proof/PPMExclusion.lean -- PPM-C with exclusion (0 sorry's)
Proof/ArithCoder.lean  -- Exact rational arithmetic coding (0 sorry's)
Proof/KolmogorovBound.lean -- Roundtrip theorem + CompressionCertificate (0 sorry's)
```

### Key Theorems (Machine-Verified)

```
K_epsSq_exact          : K_DAG(eps_sq) = 5
K_epsSq_tight          : upper = lower = 5
K_xSq_eq_5             : K_DAG(x^2 = c) = 5 (universal)
K_DAG_5_forall_small_indices : forall k v : Fin 10, K = 5
K_DAG_xSq_optimal      : x^3, x^2+x, x^2+x^2, (x^2)^2 all have K > 5
K_DAG_gap              : all non-x^2 equations have K >= 7
combined_le_sum        : combined DAG <= sum of individual DAGs
larger_axiom_lower_rest : larger axiom => lower conditional cost
measurement_tower_nontrivial : every level has K > 0
compressor_roundtrip   : decode . encode = id (PPM-C over rationals)
exclusion_zero_on_excluded : excluded symbols get zero mass
encodeStep_width       : interval width shrinks by P(symbol)
```

---

## Appendix A: Machine-Verified K-Bounds via Formalized Compression

The TLC-based complexity measure (Sections 2–9) operates on proof terms. A complementary approach yields K-bounds on arbitrary byte strings: formalize a lossless compressor in Lean 4 and prove its roundtrip property.

### A.1 Architecture

We implement PPM-C with exclusion and arithmetic coding over exact rationals in Lean 4. Four files form a chain:

PPMbayes.lean --> PPMexclusion.lean --> ArithCoder.lean --> KolmogorovBound.lean

The first file is pre-existing (I'm sorry on a measure-theoretic optimality claim, not needed for the roundtrip). The remaining three contain zero sorry's.

The central theorem:

```
theorem compressor_roundtrip (D : Nat) (data : List Byte) :
  let compressed := compressBytes (D := D) data
  decompressBytes (D := D) compressed.coder data.length = data
```

The proof proceeds by induction on the byte list. At each step, exact rational arithmetic ensures encoder and decoder compute identical intervals, the PPM-C prediction is a deterministic function of the count table (identical on both sides), and the exclusion mechanism is proved correct (`exclusion_zero_on_excluded`: excluded symbols get exactly zero mass).

## A.2 The Compression Certificate

```
structure CompressionCertificate (D : Nat) where
  original : List Byte
  compressedBits : Nat
  roundtripProof : decompressBytes (D := D)
    (compressBytes (D := D) original).coder original.length = original
  bitsPerByte : Q := compressedBits / original.length
```

This structure bundles data, compressed size, and a machine-checked proof that decompression recovers the original. It witnesses  $K(\text{original}) \leq \text{compressedBits} + |\text{decompressor}|$ . The Python implementation of the same algorithm achieves 2.16 bpb on enwik8 (see Hoekstra 2026, Appendix A of the irreversibility depth paper for the full experimental protocol), yielding  $K(\text{enwik8})/|\text{enwik8}| \leq 2.16 + O(10^{-4})$ . The empirical compression result is documented there; this appendix concerns only the formalization and the certificate.

## A.3 Why Exact Rationals

Over IEEE 754, the roundtrip holds in practice (error  $\sim 3.55 \times 10^{-15}$  per byte) but is not provable from first principles. Over the rationals, interval arithmetic is exact and the roundtrip becomes a consequence of arithmetic identities, not precision analysis.