

An End-to-End Verification of Keller’s Conjecture

Anonymous author(s)

Anonymous affiliation(s)

Abstract

In 1930, Keller conjectured that every gap-free tiling of \mathbb{R}^n by n -dimensional unit cubes must contain cubes that fully share an $(n - 1)$ -dimensional face. Keller’s conjecture holds for $n \leq 7$ and fails for $n \geq 8$. The final case, $n = 7$, was settled in 2020 using a mix of traditional and automated reasoning. The result was obtained by reducing the conjecture to a set of clique-existence problems, encoding those problems into propositional logic, breaking symmetries, and solving them with a SAT solver.

In this paper, we present an end-to-end verification in Lean 4 of Keller’s conjecture for all dimensions. First, we simplify a prior reduction of Keller’s conjecture to the clique-existence problems. We then verify an improved SAT encoding of those problems and some associated symmetry reasoning. Throughout our work, we sought to maximize the synergy between interactive and automated techniques while minimizing human proof burden. In particular, the symmetry reasoning was split between Lean and a clausal proof system, since neither was suitable on their own for verifying all the symmetry reasoning. We discuss how and why we chose to split the reasoning across these systems, based on their relative strengths and weaknesses.

2012 ACM Subject Classification Replace `ccdesc` macro with valid one

Keywords and phrases Keller’s Conjecture, the Lean theorem prover, SAT encodings, SAT solving, formal verification

Digital Object Identifier [10.4230/LIPIcs.CVIT.2016.23](https://doi.org/10.4230/LIPIcs.CVIT.2016.23)

Acknowledgements Anonymous acknowledgements

1 Introduction

In 1930, Ott-Heinrich Keller conjectured that every tiling of the n -dimensional space by unit cubes must contain adjacent cubes that fully share an $(n - 1)$ -dimensional face. Intuitively, a *tiling* T is a set of cubes that are non-overlapping, gap-free, and together cover \mathbb{R}^n . For example, all tilings of the plane by unit squares must contain two squares that share an edge (see Figure 1). Keller’s conjecture has an interesting history involving broader connections between abstract algebra and geometry; see Debroni et al. [11] for a wonderful overview.

By 2002, the conjecture was known to be true for $n \leq 6$ [11, 26] and false for $n \geq 8$ [9, 19, 29], leaving open the $n = 7$ case. Incremental results [17, 18, 32] eventually led to a proof in 2020 due to Brakensiek, Heule, Mackey, and Narváez (BHMN) [3]. By employing a mix of pen-and-paper proofs and automated reasoning tools, BHMN showed that the conjecture holds for $n = 7$, thus fully resolving the conjecture.

The mix of traditional and automated reasoning used by BHMN is becoming increasingly




Figure 1 A visual proof that Keller’s conjecture holds for $n = 2$. Either squares (a) and (b) share an edge, or the tiling must include square (c), which shares an edge with (b).

36 common. Similar proof strategies were used to resolve the Pythagorean triples problem [14],
 37 Lam’s problem [4], and the empty hexagon problem [16]. The strategy is to reduce the
 38 mathematical problem to the satisfiability of a particular propositional formula, which can
 39 then be solved by a boolean satisfiability (SAT) solver. Often, these reductions involve
 40 clever encodings, complicated symmetry reasoning, and careful case-splits to make the SAT
 41 formulas tractable to solve. However, these techniques also make the reductions prone to
 42 errors, and as a result several reductions have received formalizations in their own right,
 43 including for the Pythagorean triples [10] and empty hexagon problems [28].

44 The formalization of Keller’s conjecture has followed a similar journey, albeit a slightly
 45 longer one. BHMN reduced the conjecture to the existence of a clique in a particular family of
 46 graphs, which was then encoded into SAT. They then added symmetry-breaking constraints
 47 to the formulas to make them more tractable to solve. Some of the symmetry reasoning was
 48 handled by a mechanical proof system called *substitution redundancy* [8], but everything
 49 else was proven using pen-and-paper proofs. In 2023, Joshua Clune [6] formalized the clique
 50 reduction using the Lean 3 theorem prover, but he left open the verification of the SAT
 51 encoding and the symmetry reasoning. In this paper, we finally close this gap.

52 **Contributions.** We present the first end-to-end verification of Keller’s conjecture, which
 53 we completed using the Lean 4 theorem prover [24]. In addition to updating and significantly
 54 shortening Clune’s formalization, our novel contributions are as follows:

- 55 ■ We improved and verified BHMN’s SAT encoding of the clique problem.
- 56 ■ We verified BHMN’s symmetry reasoning.
- 57 ■ We improved BHMN’s case split to reduce SAT-solving costs.
- 58 ■ We used a verified proof checker [8] to formally check the results from the SAT solver,
 59 thus completing an end-to-end verified theorem of Keller’s conjecture.

60 Overall, our formalization shows that it is possible to effectively mix manual and automated
 61 verification in a single project without compromising correctness, and we will discuss how
 62 the different costs of interactive versus automated methods affected our formalization.

63 2 Reducing Keller to Cliques

64 In this section, we present our definition of Keller’s conjecture in Lean, and we discuss how
 65 the conjecture reduces to the existence of cliques in the so-called *Keller graphs*. While several
 66 variants of this reduction exist in the literature, we focus on the one due to BHMN [3], which
 67 was ultimately formalized by Clune [6].

68 Our formalization follows Clune’s, although we make several simplifications. In the end,
 69 our version comprises about 3,000 lines of Lean 4 code (including comments and whitespace)
 70 and represents about 1 month of work. For comparison, Clune’s formalization comprises
 71 about 15,000 lines of Lean 3 code. While part of this disparity can be attributed to the
 72 continued development of `mathlib4` [22]¹ and to improvements in Lean’s proof automation,
 73 our choice to use *Keller colorings* instead of Keller graphs (see Section 3) eliminated many
 74 lemmas and subgoals on integer arithmetic that Clune’s formalization had to grapple with.
 75 Overall, this portion of our work confirms the theme in formal verification that the right
 76 abstractions and proof automation tools can substantially reduce the costs of formalization.

77 Keller’s conjecture is usually expressed in terms of Euclidean geometry: that every tiling
 78 of \mathbb{R}^n with n -dimensional unit cubes must contain a pair of adjacent cubes that completely

¹ <https://github.com/leanprover-community/mathlib4>

```

abbrev Point (n : ℕ) := Fin n → ℝ
def UnitCube (n : ℕ) : Set (Point n) := { p | ∀ d, 0 ≤ p d ∧ p d < 1 }
def Cube (corner : Point n) : Set (Point n) := (corner + ·) '' (UnitCube n)

structure Tiling (n : ℕ) where
  corners : Set (Point n)
  covers : ∀ p : Point n, ∃! c ∈ corners, p ∈ Cube c

def Faceshare (c₁ c₂ : Point n) : Prop :=
  ∃ d, |c₁ d - c₂ d| = 1 ∧ ∀ d' ≠ d, c₁ d' = c₂ d'

def Tiling.FaceshareFree (T : Tiling n) : Prop :=
  T.corners.Pairwise (fun p₁ p₂ => ¬ Faceshare p₁ p₂)

def conjectureIn (n : ℕ) : Prop := ¬ ∃ (T : Tiling n), T.FaceshareFree

```

■ **Figure 2** The Lean definitions for Keller’s conjecture. Note that `Cube` is defined as the translation of the unit cube $[0, 1]^n$ by a `corner` point. The syntax `f '' S` means the image of `S` under `f`. The quantifier “ $\exists! x, P(x)$ ” means “there exists a unique x such that $P(x)$ holds.”

79 share an $(n - 1)$ -dimensional face. If the conjecture is false in dimension n , there exists a
80 *faceshare-free* cube tiling of \mathbb{R}^n .

81 In the literature, it is conventional to identify a tiling $T \subset \mathbb{R}^n$ with the set of minimum
82 *corners* of each cube in the tiling. A set of corners is a *tiling* if the cubes are axis-aligned,
83 non-overlapping, gap-free, and together cover \mathbb{R}^n , meaning that every point $p \in \mathbb{R}^n$ is
84 contained in a unique cube $c \in T$. To prevent double counting along (partially) shared
85 faces, the cube extending from c contains only the faces that intersect with c . In other
86 words, the cube defined by c is the half-open interval $[0, 1]^n$ translated by c , which we write
87 as $c + [0, 1]^n$. Two cubes are *facesharing* if their corners differ by a unit vector in some
88 dimension while being equal in all other dimensions. A tiling is *faceshare-free* if no pair of
89 cubes are facesharing.

90 These definitions are relatively straightforward to state in Lean; see Figure 2. Note that
91 we define points in \mathbb{R}^n as functions from `Fin n` to \mathbb{R} , rather than as vectors or lists of length n ,
92 since functions compose better in Lean. This means we access the d th coordinate of a point p
93 in Lean using function application (`p d`), rather than writing p_d or $p[d]$.

94 From here, the goal is to reduce Keller’s conjecture to the existence of a large clique in
95 a family of graphs. This reduction follows Appendix A from BHMN’s paper, which was
96 formalized by Clune. Since we repeat prior work, we only give a proof sketch.

97 A central concept of the reduction is the *core* of a tiling T . First, observe that every cube
98 with corner $c \in \mathbb{R}^n$ contains a unique integer-valued point $p \in \mathbb{Z}^n$, which we call the *index*
99 of the cube. Since the reverse is also true, i.e., every index is contained in a unique cube in T ,
100 there exists a bijection between T and \mathbb{Z}^n . The *core* of T is the set of cubes whose index is
101 in the set $I^n := \{0, 1\}^n$. Visually, these are the 2^n cubes in T that intersect the box $[0, 1]^n$,
102 with a unique cube covering each corner of the box.

103 The first half of the reduction shows that Keller’s conjecture holds for arbitrary cube
104 tilings if and only if it holds for *periodic* cube tilings. A tiling T is *periodic* if the cubes in T
105 are closed under translation by $2z$ for any $z \in \mathbb{Z}^n$. In other words, shifting the entire tiling
106 by an integer multiple of 2 units in any dimension maps the tiling back to itself. As it turns
107 out, the core of T can be used to construct a new periodic tiling T' by translating the core

108 by $2z$ for every $z \in \mathbb{Z}^n$. Visually, this means taking the cubes in T that intersect I^n and
 109 tessellating them by 2 units in every dimension. Notably, if T is faceshare-free, then so is T' .

110 The second half of the reduction shows that there exist faceshare-free periodic tilings if and
 111 only if there exist cliques of size 2^n in certain Keller graphs. A *Keller graph* $G_{n,s}$ is a graph
 112 with vertices the length- n $2s$ -ary vectors, written as $\langle 2s \rangle^n$, where $\langle 2s \rangle := \{0, 1, \dots, 2s - 1\}$.
 113 Two vertices are adjacent if they differ by s in some coordinate and also differ in some
 114 other coordinate. If K is a 2^n -sized clique for $G_{n,s}$, then a faceshare-free tiling T can be
 115 constructed from K using the points $\{\frac{v}{s} \mid v \in K\}$ as the periodic core of T . Intuitively, the
 116 s -different condition ensures that all cube corners are at least 1 unit distance apart, and the
 117 other condition ensures that possibly-facesharing corners are not actually facesharing.

118 BHMN showed that the second half of the reduction holds for the graph $G_{n,s}$ with
 119 $s = 2^{n-1}$, but since this graph is somewhat large, they used discretization results for the
 120 $n = 7$ case due to Kisielewicz and Łysakowska [17, 18, 32] to show that the reduction also
 121 holds if there exist cliques of size 2^n in any of the graphs $G_{7,s}$ for $s \in \{3, 4, 6\}$.²

122 From there, BHMN wrote a propositional formula for each $s \in \{3, 4, 6\}$, expressing
 123 the existence of a 2^7 -sized clique in $G_{7,s}$, and then they used a SAT solver to prove the
 124 non-existence of the cliques, thus showing that Keller’s conjecture was true for $n = 7$.

125 Clune formalized the reduction of Keller’s conjecture to the existence of a clique in $G_{n,s}$
 126 for $s = 2^{n-1}$, but he did not formalize Kisielewicz’s and Łysakowska’s discretization results,
 127 nor did he formalize the SAT encoding, symmetry reasoning, or SAT-solving process used
 128 to resolve the $n = 7$ case, leaving open a formalization gap. In fact, at the end of his
 129 paper, Clune remarked that closing this gap would mean either formalizing the discretization
 130 results (which would mean verifying a lot of complicated mathematics) or improving the
 131 SAT encoding so that the $G_{7,64}$ case could be tractably solved by a SAT solver.³ We chose
 132 to do the latter.

133 3 Coloring Problem


134 One of our contributions is reframing the Keller-clique problem as a coloring problem. By
 135 viewing the problem in this way, we no longer need to prove the arithmetic lemmas and
 136 subgoals regarding the s -coordinate difference between adjacent vertices in the Keller graphs,
 137 which makes the reduction much easier to verify.

138 The coloring problem formulation arises from the particular structure of Keller graphs.
 139 In 2011, Debroni et al. [11] observed that the Keller graph $G_{n,s}$ can be partitioned into
 140 2^n independent sets. See Figure 3 for an illustration. If we use the tiling construction
 141 $\{\frac{v}{s} \mid v \in K\}$ discussed previously, then the vertices of $G_{n,s}$ correspond to the points in $[0, 2)^n$
 142 arranged in a $\frac{1}{s}$ -spaced grid, and the s^n points contained in the cube extending from any
 143 *index*⁴ in $I^n := \{0, 1\}^n$ form each independent set, since no two points will be exactly 1
 144 unit apart in any dimension. For example, if $n = 2$ and $s = 2$, then the four points in the
 145 cube extending from $i = (1, 0)$, i.e., $(1, 0)$, $(\frac{3}{2}, 0)$, $(1, \frac{1}{2})$, and $(\frac{3}{2}, \frac{1}{2})$, form an independent set.
 146 Notationally, we will often write indices i as least-significant-bit-first length- n binary strings,
 147 and we will refer to their (0-indexed) d th component as i_d . For example, if $n = 7$, then the
 148 index 5 can be written as the string $i = 1010000$, and e.g. $i_2 = 1$. Two indices i and j are


² For technical reasons, it is sufficient to check only $s = 6$, but BHMN checked all three cases regardless.

³ Alternatively, one could find a SAT-less proof that there are no 2^7 -sized cliques in $G_{7,64}$, but to the best of our knowledge such a proof is not forthcoming.

⁴ The careful reader will notice that each index $i \in I^n$ is its own index, using the prior definition of an index as the unique integer-valued point in any unit cube. We re-use the term “index” on purpose.



■ **Figure 3** The Keller graph $G_{2,2}$ (left) and a partitioning of its vertices into 4 (more generally, 2^n) independent sets (right). Keller’s conjecture is false if there is a clique of size 4 (2^n), and such a clique must contain exactly one vertex from each independent set. The vertices are depicted here as “dice” with 2 (n) pips. Each pip corresponds to a different entry in that vertex’s length- n $2s$ -ary vector, with the background and dot colors corresponding to the 4 ($2s$) possible combinations for each entry. Two vertices are adjacent in the graph if one pair of corresponding pips differs only in their background color and the other pair of pips differs in any way.



■ **Figure 4** A partial Keller coloring on the indices of I^3 , depicted as the corners of a (hyper)cube. Each index receives one of s colors in each dimension. Here, only half of I^3 has been colored.


149 adjacent if they share an edge in the hypercube defined by I^n ; equivalently, they are adjacent
 150 if their binary string representations differ by only a single bit.

151 Given these independent sets, forming a clique in a Keller graph means picking one point
 152 from each index-cube in I^n . This amounts to picking a $\frac{1}{s}$ -offset for each dimension within the
 153 cube defined by each index $i \in I^n$. Equivalently, we assign one of s colors to each dimension
 154 for each index i . Figure 4 illustrates a partial coloring of the indices.

155 Putting these ideas together, we can restate the clique problem as a coloring problem.

156 ► **Definition 1** (Keller coloring). A Keller coloring in n dimensions with s colors is a mapping
 157 $C : I^n \rightarrow \langle s \rangle^n$ from indices to length- n s -ary color vectors such that the following hold:

- 158 1. **(Gap)** For all indices $i \neq j \in I^n$, there exists some dimension d for which the indices
 159 differ and the color vectors are the same: $\exists d, i_d \neq j_d \wedge C(i)_d = C(j)_d$.
- 160 2. **(No-Faceshare)** All pairs of adjacent indices i and j have pairwise different color vectors:
 161 $C(i) \neq C(j)$, or equivalently, $\exists d, C(i)_d \neq C(j)_d$.



■ **Figure 5** Examples of pairs of color vectors that do (left) and do not (right) respect the Keller coloring constraints. The background colors of the squares are the binary bits of the index, while the numbers are the colors assigned to each dimension. Top left: Gap at $d = 0$, no-faceshare at $d = 2$. Bottom left: Gap at $d \in \{0, 3\}$, not adjacent. Top right: No gap. Bottom right: Gap at $d = 3$, but adjacent and lacking a no-faceshare difference.

162 Keller colorings C are in bijection with Keller cliques $K \subset G_{n,s}$ via $K := \{v(i) \mid i \in I^n\}$,
 163 where $v(i)_d = s * i_d + C(i)_d$, and the coloring conditions correspond to the vertex-adjacency
 164 conditions in the Keller graph: The Gap condition ensures that every pair of vertices has a
 165 dimension d where their vectors differ by s , and the No-Faceshare condition ensures that the
 166 vectors of adjacent indices differ in at least one other coordinate. (Note that non-adjacent
 167 indices map to vertices with two differences no matter the coloring.)

168 We visualize (partial) colorings for certain indices as in Figure 5. Because the binary bits
 169 of the index have semantic meaning, we depict the bits with light or dark backgrounds, while
 170 the numbers in each square represent the colors assigned to each dimension.

171 Below is our definition of Keller colorings in Lean. Notably, the `KColoring` type is a
 172 structure that includes the proofs that `C` is a valid coloring, which means that the type is
 173 empty for any pair of n and s if Keller’s conjecture is true for n . We also use slightly different
 174 data types than might be expected: We use `BitVecs` rather than `Nats` or I^n to more easily
 175 access the binary representation of each index, and we use `Vectors` instead of functions from
 176 `Fin s` to `n` to increase performance when computing with colorings.

```

177 structure KColoring (n s : Nat) where
178   C : BitVec n → Vector (Fin s) n
179   gap : ∀ (i j : BitVec n), i ≠ j →
180     ∃ d : Fin n, i[d] ≠ j[d] ∧ (C i)[d] = (C j)[d]
181   no_f : ∀ (i j : BitVec n), adjacent i j →
182     ∃ d : Fin n, (C i)[d] ≠ (C j)[d]
183

```

185 4 SAT Encoding of the Keller Coloring

186 In this section, we present our CNF encoding of the Keller coloring problem. Our encoding
 187 is very similar to BHMN’s, although ours is more compact and allows for better propagation
 188 within the SAT solver. When combined with a better cube split (see Section 6), we were able
 189 to solve the coloring problem for $s = 2^{n-1}$ and therefore avoid formalizing the discretization
 190 results for the $s \in \{3, 4, 6\}$ cases.

191 We verified the encoding in Lean by reducing the Keller coloring problem to a logical
 192 specification that acted as an intermediary between the coloring problem and the encoding.
 193 We then proved that a Keller coloring exists in n dimensions with s colors if and only if the
 194 specification and the encoded propositional formula are satisfiable.

195 4.1 Boolean Satisfiability (SAT) and `trestle`

196 Given a propositional formula F , the boolean satisfiability (SAT) problem asks whether
 197 there exists a truth assignment on the boolean variables of F that satisfies it. SAT solving is
 198 an entire field in its own right; we define only the relevant concepts here. For an in-depth
 199 overview, see the Handbook of Satisfiability [2], particularly Chapters 2 and 3.

200 Almost all modern SAT solvers require the formula F to be in *conjunctive normal form*
 201 (CNF), meaning that F is written as a conjunction of clauses $F = \bigwedge C$, where each clause is
 202 a disjunction of literals $C = \bigvee \ell$. Literals ℓ are either a boolean variable x or its negation \bar{x} .
 203 Let $V(F)$ and $L(F)$ be the set of variables and literals in a formula, respectively. If F is
 204 clear from context, then we will just write V and L . Unless stated otherwise, all formulas in
 205 this paper are assumed to be in CNF.

206 A truth assignment τ is a function $\tau : L \rightarrow \{\top, \perp\}$ respecting negation. An assignment τ
 207 satisfies a CNF formula F if it satisfies at least one literal in every clause in F . Notationally,
 208 we write $\tau \models \ell$ if $\tau(\ell) = \top$, $\tau \models C$ if $\tau \models \ell$ for some $\ell \in C$, and $\tau \models F$ if $\tau \models C$ for all $C \in F$.
 209 If there exists a τ that satisfies F , then F is *satisfiable*; otherwise, it is *unsatisfiable*.

210 One of the strengths of modern SAT solvers is that they are *certifying algorithms* [23],
 211 meaning that they return answers checkable by formally-verified software. In the satisfiable
 212 case, the SAT solver returns an explicit truth assignment that can trivially be checked to
 213 satisfy all clauses in F . In the unsatisfiable case, the SAT solver emits a *proof of unsatisfiability*
 214 in a formal proof system. The de facto standard is the DRAT proof format [31]. In this paper,
 215 we use a generalization of DRAT called *substitution redundancy* (SR) [5, 8, 13] to mechanically
 216 verify symmetry breaking clauses we add to the Keller CNF encoding.

217 When we use a SAT solver, we must first *encode* the problem into a CNF formula. The
 218 choice of encoding can have a major impact on the solver’s performance, with better encodings
 219 speeding up the SAT solver by up to two orders of magnitude [27]. However, encodings are
 220 often complicated and may not obviously express the original problem, which makes verifying
 221 them an important step in the SAT solving process. Many individual SAT encodings have
 222 been verified [10, 28], and a few libraries exist for verifying SAT encodings in general [7, 12].
 223 In this formalization, we use `trestle`,⁵ a Lean 4 project for writing verified SAT encodings.
 224 The `trestle` project also contains a verified SR proof checker [8], which we use to check the
 225 SR proofs we generate.

226 Verified encodings are expressed in `trestle` using `VEncCNF`, a type that pairs an encoding `e`
 227 on a variable type ν with a proof that `e` encodes an abstract logical specification P . Roughly,
 228 an encoded formula F encodes P if all truth assignments that satisfy P also satisfy F .

```
229 abbrev PropPred ( $\nu$ ) := PropAssignment  $\nu$   $\rightarrow$  Prop
230 def VEncCNF ( $\nu$  : Type u) (P : PropPred  $\nu$ ) :=
231   { e : EncCNF  $\nu$  // e.encodesProp P }    /- In Lean, ‘//’ creates a subtype -/
232
233
```

234 In `trestle`, the encoding type `EncCNF` is a state monad that manages, among other things,
 235 the scoped declaration of auxiliary variables in a manner similar to quantifiers or lambda
 236 functions. For every new auxiliary variable v , the encoding reserves a fresh name for v in the
 237 CNF formula. This work is tedious and nontrivial, but the encoding monad hides much of it
 238 from the user, which streamlines the process of writing encodings.

239 `trestle` also provides primitives for common encoding operations. For example, the
 240 `for_all` function applies an encoding function to each element of an array and extends each
 241 sub-`VEncCNF` to one that includes the entire array:

⁵ <https://github.com/FormalSAT/trestle>. As far as we know, there is no publication for `trestle`.

```

242
243 def for_all (arr : Array  $\alpha$ ) {P :  $\alpha \rightarrow \text{PropPred } \nu$ } (f : (a :  $\alpha$ )  $\rightarrow$  VEncCNF  $\nu$  (P a))
244   : VEncCNF  $\nu$  (fun  $\tau \Rightarrow \forall a \in \text{arr}, P a \tau$ ) := ...
245

```

246 In our formalization, we make heavy use of `trestle`’s encoding primitives and `trestle`’s
 247 proof library to prove that our Keller encoding agrees with its `PropPred` specification.

248 4.2 Keller Specification and Encoding

249 We now present a CNF encoding of the Keller coloring problem, related through an inter-
 250 mediate logical *specification*. The specification and the encoding share the same boolean
 251 variables V and classes of constraints P . Crucially, the constraints in P are expressed in a
 252 form amenable to CNF encoding. Aside from the No-Faceshare constraints (see below), our
 253 encoding matches the one presented by BHMN.

254 **Variables.** We define boolean variables of the form $x_{i,d,c}$, where $i \in I^n$ is an index
 255 expressed as a length- n binary string, $d \in \langle n \rangle$ is a dimension, and $c \in \langle s \rangle$ is a color. Together,
 256 these variables represent the image of the Keller coloring C , where variable $x_{i,d,c}$ is true
 257 when index i in dimension d has color c , i.e., $\tau \models x_{i,d,c} \iff C(i)_d = c$.

258 **Constraints.** Three types of constraints comprise P : Exactly-one, Gap, and No-
 259 Faceshare, which we label as P_1 , P_2 , and P_3 , respectively. As their names suggest, they
 260 ensure that the Keller coloring C induced by the $\{x_{i,d,c}\}$ variables is well-formed and that
 261 the Gap and No-Faceshare coloring constraints are respected.

262 **Exactly-one.** The Exactly-one constraints P_1 ensure that every index i and every
 263 dimension d gets assigned a unique color c . This constraint doesn’t appear in Definition 1
 264 because it is enforced by construction, but in the propositional form we must specify that
 265 exactly one of the $x_{i,d,(.)}$ variables is true. Syntactically, we write $\exists! c, \tau \models x_{i,d,c}$, where in
 266 general $\exists! a, P(a)$ means there exists a unique a such that $P(a)$ holds.

267 To encode P_1 into a CNF, we encode “ $\exists! c$ ” with at-least-one (ALO) and at-most-one
 268 (AMO) constraints. These kinds of *cardinality constraints* are well understood in the SAT
 269 literature, and many variants of AMO constraints exist [21, 25]. For our purposes, it is
 270 sufficient to use the direct (or pairwise) AMO encoding, where binary clauses prevent every
 271 pair of relevant variables from being true at the same time. Adding ALO and AMO constraints
 272 on the same set of variables $x_{i,d,(.)}$ ensures that exactly one of them is true.

273 Formally, the specification of P_1 and its CNF encoding F_1 are:

$$274 \quad P_1(\tau) := \forall i d, \exists! c, \tau \models x_{i,d,c} \quad F_1 := \bigwedge_{i,d} \left[\left(\bigvee_c x_{i,d,c} \right) \wedge \bigwedge_{c < c'} (\bar{x}_{i,d,c} \vee \bar{x}_{i,d,c'}) \right].$$

275 **Gap.** In Definition 1, the Gap constraint says that all pairs of distinct indices i and j
 276 have a dimension d where their binary strings are different but their colors $C(i)_d$ and $C(j)_d$
 277 are the same. In terms of the x variables, this means checking that $\tau(x_{i,d,c}) = \tau(x_{j,d,c})$ for
 278 all c and some d . Formally, $P_2(\tau) := \forall i \neq j, \exists d, i_d \neq j_d \wedge \forall c, \tau(x_{i,d,c}) = \tau(x_{j,d,c})$.

279 The CNF encoding of P_2 requires a bit of cleverness. Without loss of generality, let $i < j$
 280 be two distinct indices, and let $D(i, j)$ be the set of dimensions on which i and j differ,
 281 i.e., $D(i, j) := \{d \mid i_d \neq j_d\}$. We introduce auxiliary variables $z_{i,j,d}$ (adopting the naming
 282 as in BHMN) for all $d \in D(i, j)$, where $z_{i,j,d}$ is true when i and j have the same color at
 283 dimension d . A single clause encodes that at least one $z_{i,j,d}$ is true. To encode equality, two
 284 ternary clauses enforce that if $z_{i,j,d}$ is true, then $x_{i,d,c}$ and $x_{j,d,c}$ cannot have opposite truth

285 values. Formally, the encoding F_2 is:

$$286 \quad F_2 := \bigwedge_{i < j} \left[\left(\bigvee_{d \in D(i,j)} z_{i,j,d} \right) \wedge \bigwedge_{d \in D(i,j),c} (\bar{z}_{i,j,d} \vee x_{i,d,c} \vee \bar{x}_{j,d,c}) \wedge (\bar{z}_{i,j,d} \vee \bar{x}_{i,d,c} \vee x_{j,d,c}) \right].$$

287 **No-Faceshare.** In Definition 1, the No-Faceshare constraint says that all adjacent pairs
288 of indices i and j have pairwise unequal color vectors, meaning that $C(i)_d \neq C(j)_d$ for some
289 dimension d . In terms of the x variables, this means checking that $\tau(x_{i,d,c}) \neq \tau(x_{j,d,c})$ for
290 some c and d . Formally, $P_3(\tau) := \forall i j, \text{adjacent}(i, j) \implies \exists d c, \tau(x_{i,d,c}) \neq \tau(x_{j,d,c})$.

291 The CNF encoding of P_3 uses a set of auxiliary variables similar to P_2 's. Without loss
292 of generality, let $i < j$ be two adjacent indices. Then let auxiliary variables $y_{i,j,d}$ be true
293 when i and j have *different* colors in dimension d . A single clause encodes that at least one
294 $y_{i,j,d}$ is true. To encode inequality, we use a ternary clause to disallow the case where both
295 of $x_{i,d,c}$ and $x_{j,d,c}$ are true. Formally, the encoding F_3 is:

$$296 \quad F_3 := \bigwedge_{i < j, \text{adjacent}(i,j)} \left[\left(\bigvee_d y_{i,j,d} \right) \wedge \bigwedge_{d,c} (\bar{y}_{i,j,d} \vee \bar{x}_{i,d,c} \vee \bar{x}_{j,d,c}) \right].$$


297 Our encoding of No-Faceshare differs slightly from BHMN's. We introduce variables $y_{i,j,d}$,
298 whereas BHMN introduce variables $y_{i,j,d,c}$ that perform a similar role. However, our set of
299 variables has two main advantages. First, our encoding of F_3 is more compact because it
300 uses exactly half as many clauses. Second, and more importantly, it reduces the number of
301 auxiliary variables by a factor of s . Since the formula uses $s = 64$, this means a substantial
302 reduction in the number of auxiliary variables, which in turn improves propagation speed
303 and thus performance.

304 **Trestle representation.** It is straightforward to represent the variables V and the
305 specification P in `trestle`. Here are our definitions:

```
306 inductive KellerVars (n s : Nat)
307   | x (i : BitVec n) (d : Fin n) (c : Fin s)
308
309 def KellerSpec {n s : Nat} : PropPred (KellerVars n s) := (fun τ =>
310   /- Exactly-one constraints -/
311   ( ∀ (i : BitVec n) (d : Fin n), ∃! c, τ (x i d c) ) ∧
312   /- Gap constraints -/
313   ( ∀ (i j : BitVec n), i ≠ j → ∃ d, i[d] ≠ i'[d]
314     ∧ ∀ (c : Fin s), τ (x i d c) = τ (x j d c) ) ∧
315   /- No-Faceshare constraints -/
316   ( ∀ (i j : BitVec n) (d : Fin n), i ~~~ j = BitVec.oneAt d →
317     ∃ d', d' ≠ d ∧ ∀ c, ¬ τ (x i d' c) ∨ ¬ τ (x j d' c) ) )
318
319
```

320 Expressing the constraint encodings in `trestle` is a little trickier. Whenever we declare
321 auxiliary variables $y_{i,j,d}$ or $z_{i,j,d}$, we need to prove that they have their intended semantic
322 meaning. Thankfully, `trestle` provides a primitive called `withTemps` that acts like a quantifier
323 for `VEncCNF`, much like `for_all`, but specifically for auxiliary variables. Its definition is quite
324 technical, but roughly, `withTemps` provides a way to prove that a specification on a base set
325 of variables V agrees with one that includes auxiliary variables.

326 For example, here is the `trestle` encoding of the Gap constraint on two indices. The
327 (`by ...`) block following the call to `mapProp` is the proof that the clauses in the encoding
328 satisfy the spec. Thanks to `trestle`'s proof library, this proof is only 11 lines long.



■ **Figure 6** The symmetry breaking constraints for $n = 7$ (left) and the units subsequently implied by unit propagation (right). The colored numbers indicate the fixed colorings for each dimension. The blue region is the matrix symmetry block, and the red region is the dimension 5/6 block.

```

329 def gapEncoding {n s} (i j : BitVec n) : VEncCNF (KellerVars n s)
330   (fun τ => ∃ d, i[d] ≠ j[d] ∧ ∀ c, τ (x i d c) = τ (x j d c)) :=
331   ( let dims := Array.finRange n
332     let colors := Array.finRange s
333     let D := dims.filter (fun d => i[d] ≠ j[d])
334     withTemps (Fin n) <| do /- The auxiliary variables are drawn from (Fin n) -/
335       addClause (D.map (Literal.pos (tempVar ·))) /- The ALO clause -/
336       for_all D fun d => do
337         for_all colors fun c => do
338           addClause #[Literal.neg (tempVar d), /- The two ternary clauses -/
339             Literal.pos (x i d c), Literal.neg (x j d c)]
340           addClause #[Literal.neg (tempVar d),
341             Literal.neg (x i d c), Literal.pos (x j d c)]
342         )
343     ) |>.mapProp (by ...) /- A proof that these clauses encode the PropPred spec -/
344

```

345 We were able to write the other two sub-encodings in a similar manner. Overall, it took
346 us only 150 lines of Lean code to write (and verify!) the full encoding we presented above.

347 5 Symmetry Breaking

348 The Keller coloring problem has many interesting symmetries, and as a result, so does the
349 encoding we presented in Section 4. The presence of symmetry causes the SAT solver to
350 explore redundant areas in the search space and thus time out. To make solving tractable,
351 BHMN added *symmetry-breaking constraints* to the encoded formula. Some constraints broke
352 symmetries on the level of the clique problem, while others broke symmetries in the encoded
353 formula. We call these *non-clausal* and *clausal* constraints, respectively. The non-clausal
354 constraints were justified with pen-and-paper proofs, while the clausal ones were justified
355 with mechanically-checkable proofs.

356 In our formalization, we verify the same kinds of symmetry reasoning using similar
357 methods. We verify the non-clausal constraints directly in Lean by applying insights about
358 symmetries on Keller colorings. For the clausal constraints, we use a verified substitution
359 redundancy (SR) proof checker [8] to mechanically validate the constraints we added to the
360 encoded formula. Figure 6 summarizes the set of symmetry-breaking constraints we added.

361 Overall, this portion of our work left us with two lessons. The first is that long chains of
362 symmetry reasoning are more tedious to verify in Lean than they appear on paper. This
363 is due to how each symmetry reasoning step must be shown to preserve the constraints
364 added by the previous steps. Future formalizations might be able to avoid this problem by
365 developing good abstractions for the composability of different symmetries.

$C(0)$	0	1	2	3	4	5	6	\Rightarrow	$C'(0)$	1	2	3	4	5	6	0
$C(7)$	1	1	1	0	0	0	0		$C'(67)$	1	1	0	0	0	0	1

(a) **Dimension permutation.** All dimensions shift to the left, with the first becoming the last.

$C(0)$	0	2	1	0	0	0	0	\Rightarrow	$C'(0)$	0	1	1	0	2	0	0
$C(1)$	0	1	2	0	0	0	0		$C'(1)$	0	0	2	0	2	0	0

(b) **Color permutation.** The color permutation (2 1 0) is applied to dimensions 1 and 4.

$C(42)$	0	0	0	0	0	0	0	\Rightarrow	$C'(0)$	0	0	0	0	0	0	0
$C(34)$	0	1	1	1	1	1	1		$C'(8)$	0	1	1	1	1	1	1

(c) **Index bit flip.** The index bits are flipped in dimensions 1, 3, and 5.

$C(1)$	0	0	0	0	0	0	0	\Rightarrow	$C'(1)$	0	0	0	0	0	0	0
$C(3)$	0	1	1	1	1	1	1		$C'(7)$	0	1	1	1	1	1	1

(d) **Conditional index bit flip.** The index bit is flipped in dimension 2 if the color is 1. Note that the CIBF may not preserve index adjacencies, but it does preserve the No-Faceshare coloring condition.

■ **Figure 7** Examples of symmetries being applied to colorings of various indices.

366 The second lesson is that, by carefully choosing the order in which to apply symmetries, it
 367 is possible to automatically validate a significant portion of the symmetry-breaking constraints.
 368 One of the symmetries we use, the conditional index bit flip, is not expressible in SR, and so
 369 we are forced to use Lean when we reason about it. But once we are done using it, we can
 370 (and do) mechanically validate all remaining clauses using SR. Moving this point to be as
 371 soon as possible in our proofs reduced the human proof burden.

372 5.1 Keller Coloring Symmetries

373 The Keller coloring problem is rich with symmetry. Many of these symmetries have direct
 374 geometric analogues. For example, flipping the d th bit in each index $i \in I^n$ corresponds to
 375 translating by $-e_d$ the indices with $i_d = 1$ (and their cubes) and all others by $+e_d$, where
 376 $\{e_k\}$ is the standard basis for \mathbb{R}^n . As we will see, we can use these symmetries during
 377 symmetry breaking to construct a canonical coloring with certain fixed features.

378 Our formalization uses four kinds of symmetries, illustrated in Figure 7.

- 379 a. **Dimension permutation:** Permute the order of dimensions $d \in \langle n \rangle$.
- 380 b. **Color permutation:** Permute the colors $c \in \langle s \rangle$ for a fixed dimension d .
- 381 c. **Index bit-flip:** Flip the index bit in a fixed dimension d for all indices in I^n .
- 382 d. **Conditional index bit-flip (CIBF):** Flip the index bit in a fixed dimension d , but
 383 only for indices whose color in dimension d is a given color c .

384 Transformations (a)-(c) are symmetries because they preserve all index adjacencies and
 385 dimension-wise color (in)equalities. In other words, since Keller colorings are agnostic to
 386 the ordering of the dimensions, the ordering of the colors within any dimension, and the
 387 signs of the various index bits, we are free to reorder or flip them however we wish. It is less
 388 obvious that transformation (d) is a symmetry, but it, too, is valid. Observe that a CIBF at
 389 dimension d does not affect the value of $C(i)_d$ for any $i \in I^n$, and it only flips index bits in

390 dimension d . Given these facts, it is not hard to show that the relevant color (in)equalities
 391 are preserved by the CIBF, even if the relative index adjacencies get shuffled around.

392 In Lean, we define each symmetry as a function from Keller colorings to Keller colorings.
 393 For example, here is our definition of the index bit flip. Note that `mask` is a general bit mask
 394 used to apply potentially many bit flips at once. We use a similar trick for color permutations,
 395 where the permutation function $f : \text{Fin } n \rightarrow (\text{Fin } s \simeq \text{Fin } s)$ applies a color permutation
 396 $f\ d$ to each dimension d . In Lean, `^^^` is syntax for the bitwise-XOR operation.

```
397
398 def KColoring.flip (mask : BitVec n) (K : KColoring n s) : KColoring n s where
399   C := fun i => K.C (i ^^^ mask)
400   gap := by intro i j ne
401         have := K.gap (i ^^^ mask) (j ^^^ mask) (by simpa)
402         simpa using this
403   no_f := by intro i j adj
404         exact K.no_f (i ^^^ mask) (j ^^^ mask) (by simpa [Keller.adjacent])
405
```

406 We also prove in Lean that each symmetry is, in fact, a symmetry. The proofs are
 407 extremely straightforward, and they amount to providing the inverse of each symmetry.

408 5.2 Non-Clausal Symmetry Breaking

409 Starting from the base Keller coloring problem, we add non-clausal constraints using the
 410 Keller coloring symmetries. Our goal in this section is to fix the colors in the first two rows
 411 of Figure 7 and to impose lexicographic constraints on the third row. Notably, to fix the first
 412 two rows, we assume that Keller’s conjecture holds for dimension $n - 1$. In our formalization
 413 this is not a problem, since we prove directly that the conjecture holds for $n = 2$, and then
 414 we incrementally prove that the conjecture holds for all $3 \leq n \leq 6$. Thus, we may add these
 415 symmetry-breaking constraints for any $3 \leq n \leq 7$.

416 ► **Theorem 2.** *Let $n \geq 3$, and assume Keller’s conjecture holds in dimension $n - 1$. If
 417 there exists an n -dimensional Keller coloring C , then there exists a Keller coloring C' where
 418 $C'(0) = \vec{0}$ and $C'(1) = (0, 1, \dots, \vec{0} \dots)$. In particular, if $n = 7$, then:*

$$\begin{array}{l}
 C'(0) \\
 C'(1)
 \end{array}
 \begin{array}{|c|c|c|c|c|c|c|c|}
 \hline
 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
 \hline
 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\
 \hline
 \end{array}$$

420 **Proof ([3], Lemma 2).** Let C be a valid n -dimensional Keller coloring, and consider the
 421 $(n - 1)$ -dimensional coloring \hat{C} formed by dropping the last dimension from $C(i)$ for each
 422 $i < 2^{n-1}$. Let the notation $i0$ mean appending a 0 bit to the end of index i . Then formally,
 423 we have $\hat{C}(i)_d := C(i0)_d$ for all $i \in I^{n-1}$ and $d \in \langle n - 1 \rangle$.

424 Since we assumed that Keller’s conjecture holds in dimension $n - 1$, \hat{C} cannot be a valid
 425 Keller coloring. However, observe that indices $i0$ and $j0$ respect the Gap condition in C .
 426 This means there exists a dimension $d \in \langle n \rangle$ where $i0_d \neq j0_d$ and $C(i0)_d = C(j0)_d$. Clearly,
 427 $d \neq n$, so $d \in \langle n - 1 \rangle$, which means i and j satisfy the Gap condition in \hat{C} .

428 Thus, \hat{C} must violate the No-Faceshare condition. This means there exists a pair of
 429 adjacent indices i and j where $\hat{C}(i) = \hat{C}(j)$. But since C is a valid coloring, the No-Faceshare
 430 condition in C for $i0$ and $j0$ must be satisfied by the last dimension, i.e., $C(i0)_n \neq C(j0)_n$.

431 As a result, we have found two indices $i0$ and $j0$ whose bits differ in a single dimension
 432 $d < n$ and whose color vectors differ only at dimension n . From here, we apply color
 433 permutations to set all relevant colors to 0 and 1, and then we apply bit flips and dimension
 434 permutations to map indices $(i0, j0)$ to $(0, 1)$. ◀

435 ► **Corollary 3.** *We may add the following unit clauses to the Keller CNF encoding:*

$$436 \quad \bigwedge_d x_{0,d,0} \wedge x_{1,1,1} \wedge \bigwedge_{d \neq 1} x_{1,d,0}.$$

437 Next, we apply some lexicographic sorting constraints to index 3. Some helpful notation:
 438 Let \star_k be an element of $\langle k \rangle$, let $[s:e] := \{d \mid s \leq d \leq e\}$ be the (inclusive) set of values from
 439 start s to end e , and let $J^n := \{i \in I^n \mid i = 2^d + 3, d \in [2:n]\}$ be the set of indices adjacent
 440 to 3 with an additional bit set in a single dimension in $[2:n]$.

441 ► **Theorem 4.** *Let $n \geq 3$, and let C be an n -dimensional Keller coloring with the colors fixed
 442 as in Theorem 2. Then there exists a coloring C' where $C'(3)_{[2:n]}$ is lexicographically sorted
 443 such that all positive colors appear before all 0s, and where $C'(3)_{[2:n]}$ is not lexicographically
 444 smaller than $C'(j)_{[2:n]}$ for each index $j \in J^n$ according to the ordering $0 < 1 = 2 = \dots = s$.
 445 In other words, if $d \in [2:n]$ is the greatest dimension with $C'(3)_d > 0$, then there is not an
 446 index $j \in J^n$ with $C'(j)_{d'} > 0$ for all $d' \in [2:d]$ and also $C'(j)_\delta > 0$ for some $\delta \in [d+1, n]$.*

447 **Proof.** Let $2 \leq d \leq n$ be our induction variable representing the dimension where $\bar{\star}_2$ switches
 448 from positive colors to 0s. The induction hypothesis is that $C(3)_{d'} > 0$ for every $d' \in [2:d]$.
 449 The base case of $d = n$, where index 3 has a positive color in dimensions $[2:n]$, trivially holds.

450 Now assume $d < n$. We have three cases. If there is a larger dimension $\delta > d$ with a
 451 positive color $C(3)_\delta > 0$, apply a dimension permutation swapping dimensions $d+1$ and δ ,
 452 and induct. Otherwise, there might be an index $j \in J^n$ with $C(j)_{d'} > 0$ for all $d' \in [2:d]$
 453 (the same condition as index 3) but also with a positive color $C(3)_\delta > 0$ for some $\delta > d$. If
 454 so, apply a CIBF to swap indices 3 and j . Now we are back in the first case. If such an index
 455 doesn't exist, then the conditions of the theorem hold. ◀

456 When we initially formalized this theorem, we followed BHMN's proof. However, their
 457 proof was very tedious to express in Lean, spanning 350 lines of uninteresting code. The
 458 tedium comes from proving that all prior constraints are preserved when a series of five
 459 symmetries is applied. For example, we had to show that applying a color permutation π_d
 460 mapping the positive colors $0 \neq C(3)_d \mapsto 1$ for each dimension $d \in [2:n]$ preserves the fixed
 461 colors for indices 0 and 1. These proof obligations grow as more constraints are added.

462 In the end, we found a simpler proof that spans only 200 lines of less-tedious Lean code. It
 463 helped that Theorem 4 is weaker than BHMN's: the original version of the theorem fixed the
 464 colors for index 3. We were able to weaken it by moving the fixing to the clausal symmetry
 465 breaking section, since the CIBF was only needed to lexicographically sort index 3.

466 ► **Corollary 5.** *We may add the following unit clauses to the Keller CNF encoding. Note
 467 that $(\bar{a} \vee b)$ is equivalent to $(a \Rightarrow b)$, and $(\bigwedge_A a) \Rightarrow (\bigwedge_B b)$ is equivalent to $\bigwedge_B (\bigvee_A \bar{a} \vee b)$.*

$$468 \quad \bigwedge_{d \in [2:n]} \left[(\bar{x}_{3,d,0} \vee x_{3,d+1,0}) \wedge \bigwedge_{j \in J^n} \left[\left(\bar{x}_{3,d,0} \wedge x_{3,d+1,0} \wedge \bigwedge_{d' \in [2:d]} \bar{x}_{j,d',0} \right) \Rightarrow \left(\bigwedge_{\delta \in (d:n]} x_{j,\delta,0} \right) \right] \right]$$

469 5.3 Clausal Symmetry Breaking

470 We now add symmetry-breaking clauses to the encoded CNF. These clauses are not verified
 471 in Lean directly; rather, we use recent work on the *substitution redundancy* (SR) proof
 472 system [5, 8, 13] to mechanically verify them.

473 Substitution redundancy is a *clausal proof system* that generalizes the popular RAT proof
 474 system [31]. Each step of a clausal proof either adds or deletes a clause from a CNF formula F .

475 To add a clause C , the proof checker must check that C is *redundant*, meaning that F and
 476 $F \wedge C$ are *equisatisfiable*. Here, it is sufficient to show that if F is satisfiable, so is $F \wedge C$. A
 477 proof that adds the empty clause \perp is a *proof of unsatisfiability* for F , while a proof that
 478 only adds clauses is a *proof of symmetry breaking*.

479 When adding a clause C , the proof step may optionally include a *witness* σ that helps
 480 show that C is redundant. In SR, the witness is a *substitution* $\sigma : V(F) \rightarrow \{\top, \perp\} \cup L(F)$
 481 that maps the formula’s variables to fixed truth values or to other formula literals. We apply
 482 σ to a formula F , written as $F|_\sigma$, by mapping the variables in F under σ .

483 In SR, the eponymous SR rule implies redundancy [8]. Syntactically, the SR rule is written
 484 as $F \wedge \neg C \vdash_1 (F \wedge C)|_\sigma$, where \vdash_1 denotes entailment under unit propagation. Intuitively, if
 485 an assignment τ satisfies F but not C , then σ “repairs” τ into an assignment satisfying both.

486 Many clausal proof systems, including RAT and SR, have *unhinted* (DRAT, DSR) and *hinted*
 487 (LRAT, LSR) proof formats, the main difference being that the hinted versions include unit
 488 propagation hints to make proof checking take time linear in the size of the proof. Verified
 489 proof checkers for these proof systems accept only the hinted versions.

490 In this project, we use DSR-TRIM,⁶ an unverified DSR proof checker written in C that con-
 491 verts DSR proofs into LSR proofs. We then use the verified LSR proof checker [8] implemented
 492 in `trestle` to check the hinted proofs.

493 Luckily for us, SR substitutions σ can naturally express the Keller coloring symmetries
 494 from Section 5.1. For example, suppose we want to add the unit clause $(\bar{x}_{3,2,s})$ to ensure
 495 that $C(3)_2 \neq s$. If we ever have a truth assignment that satisfies the encoded formula F
 496 but sets $x_{3,2,s}$ to true, then we may apply a color permutation swapping colors s and 1 in
 497 dimension 2. The SR substitution would then be $\sigma(x) := (x_{3,2,s} \mapsto x_{3,2,1}, x_{3,2,1} \mapsto x_{3,2,s})$,
 498 where any omitted variables are mapped to themselves. In actuality, the substitutions are
 499 more complicated than this, since we also need to remap any affected auxiliary variables
 500 $y_{3,j,2}$ and $z_{3,j,2}$, but we handle these details in our formalization.

501 For the $n = 7$ case, we generate five kinds of SR clauses. We briefly describe each.

502 **Index color bounds.** Observe that for any list of indices $L = [i_1, \dots, i_k]$, there are at
 503 most k distinct colors in any dimension d . Thus, we may apply color permutations to bound
 504 the colors for each index by its position in L , i.e., $C(i_1)_d \leq 0$, $C(i_2)_d \leq 1$, and so on.

505 We apply a similar technique to the indices 3, 7, 11, and 19 for dimensions $[2:4]$, noting
 506 that we have already fixed the color 0 for indexes 0 and 1 as in Figure 6. Thus, $C(3)_d \leq 1$,
 507 $C(7)_d \leq 2$, and so on. To accomplish this, we add the unit literals $(\bar{x}_{i,d,c})$ for $d \in [2:4]$ and
 508 for c greater than the desired bound. The SR witness is the appropriate color permutation.

509 **Index 3 fixing.** After bounding the colors for index 3, a handful of clauses become
 510 implied (RUP) by the formula. In particular, we may fix $C(3)_d = 1$ for $d \in [2:4]$.

511 **Increment sorting.** Observe that for any list of indices $L = [i_1, \dots, i_k]$ that have been
 512 color bounded, we can further enforce that the colors along any dimension d be *increment*
 513 *sorted*, meaning that each successive color is no more than one higher than the previous
 514 maximum. We add clauses that block solutions that are not increment sorted for indices 7,
 515 11, and 19 among the dimensions in $[2:5]$. The SR witnesses are color permutations.

516 **Matrix symmetry.** At this point, the 1s along the diagonal of the blue matrix are fixed,
 517 as shown in Figure 6. However, there are still matrix symmetries among the color assignments
 518 to the remaining matrix entries. As discussed in Section 6.1 of Brakensiek et al. [3], for $s \geq 4$
 519 there are 28 canonical assignments to the matrix, and for each assignment, we reorder the

⁶ <https://github.com/ccodel/dsr-trim>

520 columns of the matrix to find the lexicographically smallest matrix coloring. We add clauses
 521 to block the non-canonical assignments.

522 **Dimensions 5 and 6.** For $n = 7$, dimensions 5 and 6 are potentially still symmetric.
 523 We already lexicographically sorted index 3, so $C(3)_5 \succeq C(3)_6$. However, if these colors are
 524 equal (meaning they are both 0 or both 1), then we can swap dimensions 5 and 6 to enforce a
 525 lexicographically larger coloring in dimension 5 than in dimension 6 among the indices 7, 11,
 526 and 19. In other words, we can treat the four relevant colors in each dimension as a string
 527 (read top-down in the figure), and swap if needed to get the larger string in dimension 5.

528 **6 Solving the CNFs**

529 Finally, we solve the encoded CNF formulas with a SAT solver. To do so, we must navigate
 530 a few challenges. First, recall that the symmetry breaking for dimension n relies on Keller’s
 531 conjecture holding in dimension $n - 1$. We deal with that by proving a sequence of results
 532 from lower dimensions up to $n = 7$. Second, a significant portion of the symmetry breaking
 533 is justified with SR clauses that must be validated. Third, the formula for $n = 7$ is too
 534 hard for conventional SAT solvers, even with the improved encoding and effective symmetry
 535 breaking. To circumvent this issue, we use a common SAT-solving technique called *cube and*
 536 *conquer* [15], as was done in the original BHMN paper.

537 Starting with $n = 2$, we solve the base encoding (without any symmetry breaking) using
 538 a conventional single-threaded CDCL solver CaDiCaL, which outputs an LRAT proof. We
 539 then pass this LRAT proof to a verified LRAT checker [30], and use proof by reflection to
 540 obtain a complete proof in Lean that the conjecture holds in $n = 2$. We repeat this process
 541 for $3 \leq n \leq 5$, now including the *non-clausal* symmetry-breaking clauses, which are justified
 542 in each case by the result for the previous dimension. These results are all verified within a
 543 few seconds.

544 However, the situation is more challenging for $n = 6, 7$. In both cases, the formulas
 545 only become tractable after adding the clausal symmetry-breaking clauses. So, for both
 546 $n = 6, 7$, we use a verified tool to check that the SR proofs are correct. For $n = 6$ this check
 547 takes a few minutes, while for $n = 7$ it takes just over an hour. After adding the clausal
 548 symmetry-breaking clauses, the $n = 6$ formula can be solved by CaDiCaL quickly. But this is
 549 still insufficient for $n = 7$. We discuss this below.

550 **6.1 Cube and Conquer**

551 Cube and conquer is a technique for solving hard CNFs. A formula F is broken into many
 552 subformulas $F \wedge \alpha_i$, where the α_i are conjunctions of literals (called “cubes”) used to guide
 553 the solver [15]. As long as the cubes together form a tautology, meaning they cover the
 554 search space, solving the many subformulas is equivalent to solving F . The original BHMN
 555 result splits on all 28 remaining ways to assign the blue ‘matrix block’ and all 1378 ways to
 556 assign the red ‘dimension 5/6 columns.’ This gave a total of 38,584 cubes to check.

557 We used *proofix* [1], a new tool for automatically identifying good variables to split on,
 558 to investigate ways to improve this split. Across many runs, the tool would consistently
 559 split on just one of the variables in the matrix block. We determined that the tool was
 560 splitting between the *hardest* matrix assignment and the 27 other matrix assignments. We
 561 implemented this insight, while still splitting all 1378 cases for the red columns. As with
 562 BHMN, one cube proved substantially harder than the others. We split this cube into 16
 563 subcubes using all possible assignments to four variables suggested by *proofix* ($x_{15,5,0}$, $x_{22,0,0}$,
 564 $x_{23,6,0}$, $x_{26,0,0}$). Table 1 shows the size of the formulas and the proofs, and the time to

■ **Table 1** Statistics on solving the harder formulas. The first four columns show the dimension (n), the number of colors (s), and the number of clauses and variables of the verified encoding. The next three columns show the number of SR clauses, the time it took to add hints, and the time to validate the hinted proof (in seconds). The last three columns show the number of cubes, the total time to solve the formula, and the total time to check the proof produced by the solver (in hours).

n	s	# Cls	# Var	# SR	dsr-trim (s)	lsr-check (s)	# Cube	Solve (h)	Check (h)
6	32	617,017	25,536	1019	74.374	14.776	104	0.014	0.023
7	2	130,470	61,824	35	5.246	1.030	287	1.366	1.119
7	6	383,142	65,408	385	72.939	10.610	2771	29.332	19.578
7	64	5,657,894	117,376	2582	3991.633	370.542	2771	228.517	124.996

565 produce and validate the proofs. Compared to BHMN, our improved encoding roughly halved
 566 the solving time for $s = 6$. On subproblems, we observed that the performance improvement
 567 on $s = 64$ was more substantial. After the subformulas were all solved, we checked that the
 568 formula with the negated cubes is easy to refute, ensuring that the cubes cover the entire
 569 search space.

570 6.2 Higher Dimensions

571 For dimensions $n \geq 8$, Keller’s conjecture fails. Clune’s formalization proves that the
 572 conjecture fails for $n = 8$ by computing an explicit counterexample. This counterexample
 573 served as additional evidence that the conjecture was defined correctly in the theorem prover.
 574 However, Lean 3, the tool used in that formalization, had poor support for proof by reflection,
 575 and the $n = 8$ result took hours to check in Lean.

576 We give a formal proof that Keller’s conjecture fails for $n \geq 8$. We use Mackey’s
 577 counterexample for $n = 8$ [20], and we formally verify a program that checks that the
 578 counterexample is valid. Again using proof by reflection, we obtain a fully-verified proof
 579 that the conjecture fails for $n = 8$, and thanks to improvements between Lean 3 and Lean 4,
 580 this proof can be checked in seconds. We also prove that a counterexample in dimension n
 581 implies a counterexample in dimension $n + 1$, which is nontrivial but easy to show.

582 Thus, we formally resolve Keller’s conjecture: it holds for $n \leq 7$ and fails for $n \geq 8$.

583 7 Conclusion

584 Although we resolved Keller’s conjecture in particular, our formalization covers many standard
 585 techniques used to resolve open mathematical questions with SAT solvers. We hope that
 586 our efforts inspire further development of proof automation and tools for writing encodings,
 587 breaking symmetries, and validating SAT proofs.

588 8 Statement on AI Usage

589 We did not use LLMs or other sources of AI in the creation of our formalization, and aside
 590 from the occasional question to an LLM chatbot about Tikz styling, we did not use LLMs or
 591 other sources of AI in the writing of this paper.

592 References

- 593 1 Zachary Battleman, Joseph E. Reeves, and Marijn J. H. Heule. Problem Partitioning
594 via Proof Prefixes. In Jeremias Berg and Jakob Nordström, editors, *28th International
595 Conference on Theory and Applications of Satisfiability Testing (SAT 2025)*, volume 341
596 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 3:1–3:18, Dagstuhl,
597 Germany, 2025. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. doi:10.4230/
598 LIPIcs.SAT.2025.3.
- 599 2 Armin Biere, Marijn Heule, Hans van Maaren, and Toby Walsh, editors. *Handbook
600 of Satisfiability - Second Edition*, volume 336 of *Frontiers in Artificial Intelligence and
601 Applications*. IOS Press, 2021. doi:10.3233/FAIA336.
- 602 3 Joshua Brakensiek, Marijn Heule, John Mackey, and David Narváez. The Resolution
603 of Keller’s Conjecture. *Journal of Automated Reasoning*, 66(3):277–300, August 2022.
604 doi:10.1007/s10817-022-09623-5.
- 605 4 Curtis Bright, Kevin K. H. Cheung, Brett Stevens, Ilias Kotsireas, and Vijay Ganesh.
606 A SAT-based resolution of Lam’s Problem. *Proceedings of the AAAI Conference on
607 Artificial Intelligence*, 35(5):3669–3676, May 2021. doi:10.1609/aaai.v35i5.16483.
- 608 5 Sam Buss and Neil Thapen. DRAT and propagation redundancy proofs without new
609 variables. *Logical Methods in Computer Science*, Volume 17, Issue 2, Apr 2021. doi:
610 10.23638/LMCS-17(2:12)2021.
- 611 6 Joshua Clune. A Formalized Reduction of Keller’s Conjecture. In *Proceedings of the 12th
612 ACM SIGPLAN International Conference on Certified Programs and Proofs*, CPP 2023,
613 pages 90–101, New York, NY, USA, January 2023. Association for Computing Machinery.
614 doi:10.1145/3573105.3575669.
- 615 7 Cayden R. Codel, Jeremy Avigad, and Marijn J. H. Heule. Verified Encodings for SAT
616 Solvers. In *2023 Formal Methods in Computer-Aided Design (FMCAD)*, pages 141–151,
617 October 2023. ISSN: 2708-7824. doi:10.34727/2023/isbn.978-3-85448-060-0_22.
- 618 8 Cayden R. Codel, Marijn J. H. Heule, and Jeremy Avigad. Verified substitution redun-
619 dancy checking. In Nina Narodytska and Philipp Rümmer, editors, *Proceedings of the
620 24th Conference on Formal Methods in Computer-Aided Design - FMCAD 2024*, volume 5
621 of *Formal Methods in Computer-Aided Design*, pages 186–196, Vienna, Austria, 10 2024.
622 TU Wien Academic Press. doi:10.34727/2024/isbn.978-3-85448-065-5_24.
- 623 9 K. Corrádi and S. Szabó. A combinatorial approach for Keller’s conjecture. *Periodica
624 Mathematica Hungarica*, 21(2):95–100, June 1990. doi:10.1007/BF01946848.
- 625 10 Luís Cruz-Filipe, João Marques-Silva, and Peter Schneider-Kamp. Formally verifying the
626 solution to the Boolean Pythagorean triples problem. *J. Autom. Reason.*, 63(3):695–722,
627 2019.
- 628 11 Jennifer Debroni, John D. Eblen, Michael A. Langston, Wendy Myrvold, Peter Shor, and
629 Dinesh Weerapurage. A complete resolution of the Keller maximum clique problem. In
630 *Proceedings of the Twenty-Second Annual ACM-SIAM Symposium on Discrete Algorithms*,
631 pages 129–135. Society for Industrial and Applied Mathematics, January 2011. doi:
632 10.1137/1.9781611973082.11.
- 633 12 Sofia Giljegård and Johan Wennerbreck. Puzzle solving with proof. Master’s thesis,
634 Chalmers University of Technology, University of Gothenburg, 2021.
- 635 13 Stephan Gocht and Jakob Nordström. Certifying parity reasoning efficiently using
636 pseudo-boolean proofs. In *Thirty-Fifth AAAI Conference on Artificial Intelligence*,
637 *AAAI 2021*, *Thirty-Third Conference on Innovative Applications of Artificial Intelligence*,
638 *IAAI 2021*, *The Eleventh Symposium on Educational Advances in Artificial Intelligence*,

- 639 *EAAI 2021, Virtual Event, February 2-9, 2021*, pages 3768–3777. AAAI Press, 2021.
640 [doi:10.1609/AAAI.V35I5.16494](https://doi.org/10.1609/AAAI.V35I5.16494).
- 641 14 Marijn J. H. Heule, Oliver Kullmann, and Victor W. Marek. Solving and verifying the
642 Boolean Pythagorean triples problem via cube-and-conquer. In Nadia Creignou and
643 Daniel Le Berre, editors, *Theory and Applications of Satisfiability Testing – SAT 2016*,
644 pages 228–245, Cham, 2016. Springer International Publishing.
- 645 15 Marijn J. H. Heule, Oliver Kullmann, Siert Wieringa, and Armin Biere. Cube and
646 conquer: Guiding CDCL SAT solvers by lookaheads. In Kerstin Eder, João Lourenço,
647 and Onn Shehory, editors, *Hardware and Software: Verification and Testing*, pages 50–65,
648 Berlin, Heidelberg, 2012. Springer Berlin Heidelberg.
- 649 16 Marijn J. H. Heule and Manfred Scheucher. Happy ending: An empty hexagon in every
650 set of 30 points. In *Tools and Algorithms for the Construction and Analysis of Systems:
651 30th International Conference, TACAS 2024*, pages 61–80, Berlin, Heidelberg, 2024.
652 Springer-Verlag. [doi:10.1007/978-3-031-57246-3_5](https://doi.org/10.1007/978-3-031-57246-3_5).
- 653 17 Andrzej P. Kisielewicz. Rigid polyboxes and Keller’s conjecture, December 2014.
654 arXiv:1304.1639 [math]. [doi:10.48550/arXiv.1304.1639](https://doi.org/10.48550/arXiv.1304.1639).
- 655 18 Andrzej P. Kisielewicz and Magdalena Lysakowska. On Keller’s Conjecture in Dimension
656 Seven. *The Electronic Journal of Combinatorics*, 22(1):P1.16, January 2015. [doi:
657 10.37236/4153](https://doi.org/10.37236/4153).
- 658 19 Jeffrey C. Lagarias and Peter W. Shor. Keller’s cube-tiling conjecture is false in high
659 dimensions. 1992. Publisher: arXiv Version Number: 1. [doi:10.48550/ARXIV.MATH/
660 9210222](https://doi.org/10.48550/ARXIV.MATH/9210222).
- 661 20 Mackey. A Cube Tiling of Dimension Eight with No Facesharing. *Discrete & Computa-
662 tional Geometry*, 28(2):275–279, August 2002. [doi:10.1007/s00454-002-2801-9](https://doi.org/10.1007/s00454-002-2801-9).
- 663 21 Joao Marques-Silva and Inês Lynce. Towards robust CNF encodings of cardinality con-
664 straints. In Christian Bessière, editor, *Principles and Practice of Constraint Programming
665 – CP 2007*, pages 483–497, Berlin, Heidelberg, 2007. Springer Berlin Heidelberg.
- 666 22 The Mathlib Community. The Lean mathematical library. In *Proceedings of the 9th
667 ACM SIGPLAN International Conference on Certified Programs and Proofs, POPL ’20*,
668 Jan 2020. [doi:10.1145/3372885.3373824](https://doi.org/10.1145/3372885.3373824).
- 669 23 Ross M. McConnell, Kurt Mehlhorn, Stefan Näher, and Pascal Schweitzer. Certifying
670 algorithms. *Comput. Sci. Rev.*, 5(2):119–161, 2011. [doi:10.1016/J.COSREV.2010.09.
671 009](https://doi.org/10.1016/J.COSREV.2010.09.009).
- 672 24 Leonardo de Moura and Sebastian Ullrich. The Lean 4 theorem prover and programming
673 language. In *Automated Deduction – CADE 28*, pages 625–635, 2021.
- 674 25 Van-Hau Nguyen, Van-Quyet Nguyen, Kyungbaek Kim, and Pedro Barahona. Empirical
675 study on SAT-encodings of the at-most-one constraint. In *The 9th International Confer-
676 ence on Smart Media and Applications, SMA 2020*, pages 470–475, New York, NY, USA,
677 2021. Association for Computing Machinery. [doi:10.1145/3426020.3426170](https://doi.org/10.1145/3426020.3426170).
- 678 26 Oskar Perron. Über lückenlose Ausfüllung des n -dimensionalen Raumes durch kongru-
679 ente Würfel. *Mathematische Zeitschrift*, 46(1):1–26, December 1940. [doi:10.1007/
680 BF01181421](https://doi.org/10.1007/BF01181421).
- 681 27 Bernardo Subercaseaux and Marijn J. H. Heule. The packing chromatic number of the
682 infinite square grid is 15. In Sriram Sankaranarayanan and Natasha Sharygina, editors,
683 *Tools and Algorithms for the Construction and Analysis of Systems (TACAS)*, pages
684 389–406, Cham, 2023. Springer Nature Switzerland.
- 685 28 Bernardo Subercaseaux, Wojciech Nawrocki, James Gallicchio, Cayden Codel, Mario
686 Carneiro, and Marijn J. H. Heule. Formal Verification of the Empty Hexagon Number.

- 687 In *15th International Conference on Interactive Theorem Proving (ITP 2024)*, volume
688 309 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 35:1–35:19, 2024.
- 689 **29** S. Szabó. A reduction of Keller’s conjecture. *Periodica Mathematica Hungarica*, 17(4):265–
690 277, December 1986. doi:[10.1007/BF01848388](https://doi.org/10.1007/BF01848388).
- 691 **30** Yong Kiam Tan, Marijn J. H. Heule, and Magnus O. Myreen. Verified propagation redun-
692 dancy and compositional unsat checking in cakeml. *International Journal on Software
693 Tools for Technology Transfer*, 25(2):167–184, 2023. doi:[10.1007/s10009-022-00690-y](https://doi.org/10.1007/s10009-022-00690-y).
- 694 **31** Nathan Wetzler, Marijn J. H. Heule, and Warren A. Hunt. DRAT-trim: Efficient checking
695 and trimming using expressive clausal proofs. In Carsten Sinz and Uwe Egly, editors,
696 *Theory and Applications of Satisfiability Testing - SAT 2014*, volume 8561, pages 422–429.
697 Springer International Publishing, 2014. doi:[10.1007/978-3-319-09284-3_31](https://doi.org/10.1007/978-3-319-09284-3_31).
- 698 **32** Magdalena Łysakowska. Extended Keller graph and its properties. *Quaestiones Mathe-*
699 *maticae*, 42(4):551–560, April 2019. doi:[10.2989/16073606.2018.1462865](https://doi.org/10.2989/16073606.2018.1462865).