COUNTING OPTIMAL DISCRETE MORSE MATCHINGS: THE SIX-VERTEX PROJECTIVE PLANE HAS EXACTLY 428 400

KOREA SUPERINTELLIGENCE LABS

Abstract. A matching W on the Hasse diagram of a finite simplicial complex K is a *discrete Morse matching* when reversing its edges leaves no directed cycle; the unmatched simplices are *critical*, and W is *optimal* when their number is minimal. How many optimal matchings does a complex have? For the simplex the answer is a published sequence with four terms, $f(1), \ldots, f(4) =$ 2, 9, 256, 380 125, the last of them a cluster computation; at $n = 5$ even the f-vector that would contain the next term is called computationally infeasible. Every count we could find in the literature is a count for a simplex or for the boundary of one. We compute an exact count for a complex that is neither: the 6-vertex triangulation $\mathbb{RP}^2_6$ of the real projective plane has exactly 428 400 optimal discrete Morse matchings, and every one of them has Morse vector $(1, 1, 1)$. In the notation of Zheng's Morse ensemble polynomial this says $p_{\mathbb{F}_2}(\mathbb{RP}^2_6) = 428\,400$ and $p_{\mathbb{Q}}(\mathbb{RP}^2_6) = 0$: a complex with optimal matchings in abundance and no rationally perfect one, which is where Zheng's "perfect" and Scoville's "optimal" part company. It is a new data point above dimension one for Zheng's second open problem, "Perfect Morse counts in higher dimensions", whose only such datum is $[z_0 z_2]\, \mathrm{ME}_{\partial\Delta^3} = 256$. The proof is an exhaustive enumeration of the 1 698 480 matchings of size 14 and the 70 144 matchings of size 15 on the 31-cell Hasse diagram, testing acyclicity at every leaf and pruning by acyclicity nowhere. It is machine-checked in Lean 4, as are reproductions of every published number the paper leans on: the Chari–Joswig values $f(1), f(2), f(3)$, Scoville's two f-vectors and his $f(4) = 380\,125$, three of Zheng's Morse ensemble polynomials, Zheng's identity $p(G) = n\,\tau(G)$ for connected graphs, and the optimum values $c = 3$ that Joswig–Pfetsch obtain by branch and cut for the projective plane and for the dunce hat. Counts for the 7-vertex torus and the dunce hat, 59 523 534 and 63 617 968, are reported as measurements outside the formal development and are labelled as such wherever they appear.

1. Introduction

The objects. Let K be a finite abstract simplicial complex. Its *Hasse diagram* $H(K)$ is the directed graph whose vertices are the simplices of K, with an edge $\sigma \to \tau$ whenever $\sigma \subsetneq \tau$ and $\dim \tau = \dim \sigma + 1$. A *matching* on $H(K)$ is a set W of edges of $H(K)$, no two of which share a simplex. Writing $H_W(K)$ for the digraph obtained from $H(K)$ by reversing every edge of W, the matching W is a *discrete Morse matching* — equivalently an *acyclic matching*, equivalently a gradient vector field in the sense of Forman [3] — when $H_W(K)$ has no directed cycle. A simplex not covered by W is *critical*; the *Morse vector* of W is the sequence $c(W) = (c_0(W), \ldots, c_{\dim K}(W))$ counting the critical simplices by dimension, and W is *optimal* when $\sum_i c_i(W)$ is minimal, equivalently when $|W|$ is maximal. These are the definitions of [7, §2], in Zheng's notation c_i for the critical counts rather than Scoville's m_i.

Two counting functions attach to this. Chari and Joswig [2] introduced the complex $M(K)$ of discrete Morse matchings, a simplicial complex whose vertices are the edges of $H(K)$ and whose faces are the acyclic matchings, and they write $f(n)$ for the number of optimal matchings on the n-simplex Δ^n — the number of top-dimensional faces of $M(\Delta^n)$. Zheng [9] attaches to K the *Morse ensemble polynomial*

$$\mathrm{ME}_K(z_0, \ldots, z_d) \;=\; \sum_{W \text{ acyclic}} \prod_{i=0}^{d} z_i^{c_i(W)},$$

which records the whole distribution of Morse vectors, and calls a matching *perfect* when $c_i(W) = \beta_i(K)$ for every i; the number of these over a field k is the *perfect coefficient* $p_k(K) = [z^{\beta(K;k)}]\,\mathrm{ME}_K$ of [9, §5.5].

Where the published data stop. The values of f are known for $n \leq 4$. Chari and Joswig [2] record $f(1) = 2$, $f(2) = 9$, $f(3) = 256$, and Scoville opens [7, §3] with

> "We computed that $f(4) = 380{,}125$ by exhaustive enumeration of all maximum acyclic matchings on the Hasse diagram of Δ^4, carried out on a computer cluster."

The same section states that no closed formula for $f(n)$ is known, and [7, §5.2] that for $n = 5$ even the f-vector of $M(\Delta^5)$ is computationally infeasible. Zheng [9] computes ME_K completely for graphs, where the perfect coefficient of a connected graph on n vertices is $p(G) = n\,\tau(G)$ with τ the number of spanning trees, and closes with a list of open problems. The second of them, "Perfect Morse counts in higher dimensions", asks for the higher-dimensional counterpart of that identity: whether the perfect coefficients of complexes of dimension at least two are governed by cellular spanning-tree enumerators. Its motivation is a single datum above dimension one, $[z_0 z_2]\,\mathrm{ME}_{\partial\Delta^3} = 256 = 4 \cdot 4 \cdot \tau(K_4)$, and that is the only perfect coefficient of a complex of dimension ≥ 2 stated anywhere in [9].

So the published counts we could find are counts for Δ^n and $\partial\Delta^n$, and they stop at $n = 4$. For every other complex the literature we searched reports the *optimum value* and not the number of optima: Joswig and Pfetsch [5] solve the optimal-matching problem by branch and cut over a library of triangulations, and their tables give, for each complex, the number of cells, the number of Hasse edges, a Betti lower bound and the optimum number c of critical faces — one optimal matching, never a count. Their table contains the two complexes we take up here under the names `projective` and `dunce`. Benedetti and Lutz's random discrete Morse theory [1] reports outcomes as hit frequencies, "X out of 10 000 runs", which is again not a count.

Remark 1.1 (what the negative statement rests on). "Every count we could find" is a report of a search, not a theorem. We searched the OEIS for $2, 9, 256, 380\,125$ and for each of the counts below, and by keyword; an 86 GB local full-text mirror of the mathematics arXiv, for the digit strings and by keyword; the arXiv API for `all:"Morse matchings"` and `all:"optimal Morse matching"`; and the citation list of [2] in OpenAlex, read by title with the plausible entries read in full. Nothing found counts optimal matchings on a complex other than Δ^n or $\partial\Delta^n$. The statements about [5] and [1] above are not part of that negative: both papers were read first-hand. A referee may know a source we did not reach, and that would cost this paper only its claim to be first, not any of its numbers.

Question 1.2. How many optimal discrete Morse matchings does a complex that is not a simplex or the boundary of one have?

What is settled here. Let $\mathbb{RP}^2_6$ be the 6-vertex triangulation of the real projective plane, with facets

$$012,\ 013,\ 024,\ 035,\ 045,\ 125,\ 134,\ 145,\ 234,\ 235 \tag{1}$$

on the vertex set $\{0, \ldots, 5\}$ (Proposition 3.2 identifies it as $\mathbb{RP}^2$ from the facet list alone). Its f-vector is $(6, 15, 10)$, so it has 31 simplices, and its Hasse diagram has 60 edges — the pair $(n, m) = (31, 60)$ of the `projective` row of [5, Tables 1–2], which pins our complex to theirs by more than a name.

Theorem 1.3. *On the Hasse diagram of $\mathbb{RP}^2_6$ no matching of size 15 is acyclic, and exactly 428 400 of the matchings of size 14 are acyclic. Hence the maximum size of a discrete Morse matching on $\mathbb{RP}^2_6$ is 14, the minimum number of critical simplices is 3, and*

$$\#\{\textit{optimal discrete Morse matchings of } \mathbb{RP}^2_6\} \;=\; 428\,400.$$

Theorem 1.4. *Every optimal discrete Morse matching of* $\mathbb{RP}^2_6$ *has Morse vector* $(1,1,1)$*: of the three Morse vectors with three critical simplices available on the* f*-vector* $(6,15,10)$*, namely* $(1,1,1)$*,* $(2,1,0)$ *and* $(0,1,2)$*, the last two are realised by no acyclic matching. Consequently*

$$p_{\mathbb{F}_2}(\mathbb{RP}^2_6) = 428\,400, \qquad p_{\mathbb{Q}}(\mathbb{RP}^2_6) = 0,$$

and the optimal matchings of $\mathbb{RP}^2_6$ *are exactly its* $\mathbb{F}_2$*-perfect ones, while it has no rationally perfect discrete Morse matching at all.*

Of these the number 428 400 is what is new. The *minimum* it attains is not: three is the value c reported for `projective` in [5, Tables 1–2], and it is also what the weak Morse inequalities over $\mathbb{F}_2$ give, since $\beta(\mathbb{RP}^2;\mathbb{F}_2) = (1,1,1)$. We re-derive it below by exhaustion rather than import it, because the count and the minimum come out of the same search and it would be strange to prove one and cite the other; but the re-derivation is a control, not a result.

Theorem 1.4 is where the two words in circulation part company. On Δ^n and $\partial\Delta^n$ — the only complexes for which we found published counts — "perfect" in the sense of [2, 9] and "optimal" in the sense of [7] agree, and over any field. On $\mathbb{RP}^2_6$ they do not: over $\mathbb{F}_2$ the two notions still coincide and count 428 400, while over $\mathbb{Q}$ the Betti vector is $(1,0,0)$ and no acyclic matching realises it, so the perfect coefficient vanishes although optimal matchings exist in abundance. A search of the literature for the word "perfect" alone therefore misses exactly the complexes this paper is about. We count *optimal* matchings, in Scoville's sense, and report the Morse vector separately, so that both readings are covered.

Two further complexes, measured rather than proved. The same enumeration was run on two more complexes with the same f-vector shape: the 7-vertex Möbius triangulation T^2_7 of the torus and Hachimori's 8-vertex dunce hat [4], the complex called `dunce` in [5]. For both, the minimum number of critical simplices is established here by machine-checked exhaustion — four for the torus (Proposition 5.1), three for the dunce hat (Proposition 5.2, which re-derives Zeeman's non-collapsibility [8] and the value $c = 3$ of [5] from the definition). The two *counts*, 59 523 534 and 63 617 968, are not part of the formal development: they are reported in Remark 5.3 as measurements produced by independent enumerators outside it, and Section 6 gives the measured price of promoting them. We keep that boundary visible everywhere it matters.

Organisation. Section 2 fixes definitions and the elementary arithmetic of Morse vectors on a 2-complex. Section 3 describes the enumerator and the controls it passes before it is asked for a new number, including machine-checked reproductions of every published value this paper stands on. Section 4 proves Theorems 1.3 and 1.4. Section 5 treats the torus and the dunce hat. Section 6 records costs and what was deliberately not run. Section 7 is a note on the source paper. Section 8 describes the machine verification.

2. Preliminaries

Complexes, Hasse diagrams, matchings. All complexes are finite and abstract: K is a set of nonempty finite subsets of $\{0,1,\ldots,n_v-1\}$ closed under passing to nonempty subsets, and we call its members *cells* or *simplices.* We specify a complex by its facets and mean the complex they generate. The Hasse diagram $H(K)$, matchings, the modified diagram $H_W(K)$, acyclic matchings, critical cells, the Morse vector $c(W)$ and optimality are as in Section 1.

Two conventions are worth naming. First, a matching is a set of *cover pairs* (σ,τ) with $\sigma \subsetneq \tau$ and $\dim\tau = \dim\sigma + 1$, and no cell may lie in two pairs; a matching of size $|W| = k$ on a complex with N cells leaves exactly $N - 2k$ critical cells, so

$$\textstyle\sum_i c_i(W) \;=\; N - 2|W| \;\equiv\; N \pmod 2. \tag{2}$$

Maximising $|W|$ over acyclic matchings and minimising the number of critical cells are therefore the same problem. Second, the direction in which one reverses is immaterial: a finite digraph has a directed cycle if and only if its reverse does, so orienting $H(K)$ downwards and reversing matched edges upwards — the convention of the computations below — gives the same acyclic matchings as Scoville's upward orientation.

We use two classical facts about acyclic matchings and cite them rather than reprove them: Forman's theorem that K is homotopy equivalent to a CW complex with $c_i(W)$ cells in each dimension i [3], and hence the weak Morse inequalities $c_i(W) \geq \beta_i(K;k)$ over every field k; and the fact that a subset of an acyclic matching is acyclic, which makes $M(K)$ a simplicial complex [2]. The second is used only to speak of the f-vector of $M(K)$ as the f-vector of a complex; *no computation below assumes it*, a point we return to in Section 3.

Morse vectors on a 2-complex. The following bookkeeping is elementary but is used repeatedly, and it is what makes the case analysis in Theorem 1.4 finite.

Lemma 2.1. *Let K be a 2-dimensional complex with f-vector (f_0, f_1, f_2) and let W be a matching on $H(K)$ with Morse vector (c_0, c_1, c_2). Then W consists of $f_0 - c_0$ vertex–edge pairs and $f_2 - c_2$ edge–triangle pairs, and*

$$c_1 \;=\; c_0 + c_2 - (f_0 - f_1 + f_2) \;=\; c_0 + c_2 - \chi(K).$$

In particular, if $c_0 + c_1 + c_2 = s$ then $c_1 = (s - \chi(K))/2$ and $c_0 + c_2 = (s + \chi(K))/2$.

Proof. Every pair of W is a cover pair, hence either vertex–edge or edge–triangle. Counting matched vertices gives $f_0 - c_0$ pairs of the first kind, counting matched triangles gives $f_2 - c_2$ of the second, and counting matched edges gives $f_1 - c_1 = (f_0 - c_0) + (f_2 - c_2)$, which rearranges to the display. Adding $c_0 + c_2 = c_1 + \chi(K)$ to $c_0 + c_1 + c_2 = s$ gives the last sentence. □

For $\mathbb{RP}^2_6$, with $(f_0, f_1, f_2) = (6, 15, 10)$ and $\chi = 1$, a matching with three critical cells has $c_1 = 1$ and $c_0 + c_2 = 2$, so its Morse vector is one of

$$(1, 1, 1), \qquad (2, 1, 0), \qquad (0, 1, 2), \tag{3}$$

and these three are what Theorem 1.4 has to separate. For T^2_7, with $(7, 21, 14)$ and $\chi = 0$, a matching with four critical cells has Morse vector $(1, 2, 1)$, $(2, 2, 0)$ or $(0, 2, 2)$; for the dunce hat, with $(8, 24, 17)$ and $\chi = 1$, a matching with three critical cells has Morse vector $(1, 1, 1)$, $(2, 1, 0)$ or $(0, 1, 2)$.

3. The enumerator, and the controls it passes first

What is computed, and what is not assumed. Every number in this paper is produced by one recursion. A complex is built from its facet list as flat integer data: the cells are all nonempty faces of the facets, indexed in increasing order of their vertex bitmask, so that a cell's covers always have larger index than the cell; from that indexing the covering lists, the coboundary lists and the successor lists of the modified diagram are tabulated once. Nothing about the diagram is transcribed by hand — the only hand-entered data anywhere are the facet lists, and Propositions 3.1 and 3.2 pin those to the complexes they are supposed to triangulate.

The search then walks the cells in index order. At each cell not yet used it branches: leave the cell critical, or match it with one of its not-yet-used covers. The *only* pruning is the cardinality bound "at most $N - 2k$ cells may be left critical if k pairs are to be reached", or in the profile-constrained variant the per-dimension budget of Lemma 2.1; both are facts about counting alone. Acyclicity is tested only at the leaves, by Kahn's algorithm — repeatedly delete a source from $H_W(K)$ and check that the digraph is exhausted.

That design is deliberate and is the reason the counts below are what they claim to be. A backtracking search over acyclic matchings normally prunes as soon as the partial matching has a cycle, which is licensed by the heredity of acyclicity under passing to subsets. Our search never does this, so *no property of acyclicity under restriction of a matching is used anywhere in the formal development*. The price is a larger search tree; the gain is that the statement proved is exactly "the number of acyclic matchings of the prescribed size (or profile) is . . . ", with no lemma interposed.

Controls on the instrument. Before the enumerator is asked for a number nobody has published, it is checked against everything that can be checked.

Proposition 3.1. *Let A be Kahn's source-deletion test for acyclicity of $H_W(K)$ and let B be the test that iterates the successor relation of $H_W(K)$ to its transitive closure and asks whether some cell reaches itself. Then A and B return the same verdict on* every *matching, acyclic or not, of the 3-cycle, of Δ^2, of Δ^3 and of $\partial\Delta^3$. Moreover the test is not vacuous: on the 3-cycle with vertices v_0, v_1, v_2, the matching*

$$\{\,(v_0,\, v_0v_1),\quad (v_1,\, v_1v_2),\quad (v_2,\, v_0v_2)\,\}$$

— a legal matching of the Hasse diagram whose modified diagram is a closed V-path — is rejected, while the submatching $\{(v_0,\, v_0v_1),\, (v_1,\, v_1v_2)\}$ and the empty matching are accepted. Finally the Hasse data are as stated: Δ^4 has 31 cells and 75 Hasse edges, $\partial\Delta^4$ has 30 and 70, $\mathbb{RP}^2_6$ has 31 and 60, and T^2_7 has 42 and 84.

The two tests share no code beyond the successor table: one is a linear-time source deletion, the other a cubic-time closure. The $(31, 60)$ and, below, $(49, 99)$ shape data are what identify our $\mathbb{RP}^2_6$ and our dunce hat with the `projective` and `dunce` rows of [5, Tables 1–2].

Proposition 3.2. *The facet list (1) defines a closed surface with f-vector $(6, 15, 10)$ and $\chi = 1$ which is non-orientable, hence is a triangulation of $\mathbb{RP}^2$; the facet list*

(4) 012, 013, 024, 035, 046, 056, 125, 136, 145, 146, 234, 236, 256, 345

defines a closed surface with f-vector $(7, 21, 14)$, $\chi = 0$ and orientable, hence a triangulation of the torus; and $\partial\Delta^3$ is a closed surface with f-vector $(4, 6, 4)$ and $\chi = 2$, hence a triangulation of S^2. The test is not vacuous: neither Δ^2 nor the 3-cycle is a closed surface.

Here "closed surface" is the combinatorial condition — the complex is 2-dimensional, every edge lies in exactly two triangles, every vertex link is a single cycle, and the complex is connected — and orientability is decided by propagating a coherent orientation across the dual graph. The identification of the underlying space then uses the classification of closed surfaces, which is the one piece of topology we cite rather than compute.

Remark 3.3. Outside the formal development, a separate program enumerated every 6-vertex closed surface with 10 triangles and every 7-vertex closed surface with 14 triangles, up to isomorphism: there is exactly one of each. So "the" 6-vertex $\mathbb{RP}^2$ and "the" 7-vertex torus are well defined, and the counts below are invariants of the complexes rather than of a labelling of their vertices. This census is a computation outside the formal development and nothing above or below depends on it.

Published values, reproduced. The next five propositions reproduce, from the definitions and by the same enumerator, every published number this paper stands on. None of them is new; they are stated because a count of 428 400 is worth exactly as much as the machine that produced it, and because Proposition 3.8 in particular is a genuine agreement of two independent algorithms with a published closed form.

Proposition 3.4 (Chari–Joswig)**.** *$f(1) = 2$, $f(2) = 9$ and $f(3) = 256$: the numbers of optimal discrete Morse matchings on Δ^1, Δ^2 and Δ^3 are the values recorded in* [2].

Proposition 3.5 (Scoville's f-vectors)**.** *The complexes of discrete Morse matchings of Δ^3 and of $\partial\Delta^3$ have f-vectors*

$$f\big(M(\Delta^3)\big) = (28, 300, 1544, 3932, 4632, 2128, 256), \qquad f\big(M(\partial\Delta^3)\big) = (24, 216, 896, 1692, 1248, 256),$$

as in [7, Table 1]*. Here f_i is the number of acyclic matchings with $i+1$ pairs, computed by a full census of the acyclic matchings of every size on the two Hasse diagrams.*

Proposition 3.6 (Scoville's $f(4)$)**.** $f(4) = 380\,125$*: there are exactly* $380\,125$ *acyclic matchings with* 15 *pairs on the* 31*-cell Hasse diagram of* Δ^4*. Moreover* $\partial\Delta^4$*, which has* 30 *cells, carries exactly* $380\,125$ *acyclic matchings with* 14 *pairs and none with* 15*; the first of these is the case* $n = 4$ *of* [7, Proposition 9]*, a bijection between the top-dimensional faces of* $M(\Delta^n)$ *and those of* $M(\partial\Delta^n)$ *for* $n \geq 2$*, and the second says that* S^3 *has no discrete Morse matching with no critical cell.*

Proposition 3.7 (Zheng's Morse ensemble polynomials)**.** *For a graph* G *on* n *vertices with* m *edges, every matched pair is a vertex–edge pair, so a matching with* s *pairs has Morse vector* $(n-s, m-s)$ *and* ME_G *is determined by the census of acyclic matchings by size. With that,*

$$\begin{aligned} \mathrm{ME}_{K_4} &= 64z_0z_1^3 + 48z_0^2z_1^4 + 12z_0^3z_1^5 + z_0^4z_1^6, \\ \mathrm{ME}_{G_1} = \mathrm{ME}_{G_2} &= 72z_0z_1^2 + 192z_0^2z_1^3 + 176z_0^3z_1^4 + 73z_0^4z_1^5 + 14z_0^5z_1^6 + z_0^6z_1^7, \\ [z_0z_2]\ \mathrm{ME}_{\partial\Delta^3} &= 256, \end{aligned}$$

where K_4 *is taken as a one-dimensional complex — "*$\mathrm{Ind}(U_{2,4})$ *is the one-dimensional complex* K_4*" is how* [9, §6] *puts it, and the first line is stated there for* $\mathrm{Ind}(U_{2,4})$ *— and where* G_1, G_2 *are the graphs on* $\{0, \ldots, 5\}$ *with edge sets*

$$E(G_1) = \{01,\ 02,\ 03,\ 05,\ 14,\ 23,\ 45\}, \qquad E(G_2) = \{01,\ 02,\ 03,\ 14,\ 15,\ 24,\ 25\},$$

with degree sequences $(4, 2, 2, 2, 2, 2)$ *and* $(3, 3, 3, 2, 2, 1)$*. Zheng exhibits them in* [9, Example 4.6] *as a non-isomorphic pair with the same Laplacian spectrum, and all three values are his.*

Proposition 3.8 (Zheng's identity, against an independent spanning-tree count)**.** *For a connected graph* G *on* n *vertices with* m *edges, a perfect matching in Zheng's sense is an acyclic matching with Morse vector* $(1, m-n+1)$*, and* [9, Theorem 4.9] *proves* $p(G) = n\,\tau(G)$*. For* C_3*,* C_5*,* K_4*,* K_5*,* $K_{3,3}$ *and Zheng's* G_1 *the left-hand side, computed by the matching search, and the right-hand side, computed by an independent brute-force enumeration of spanning trees, agree, both being*

$$9, \quad 25, \quad 64, \quad 625, \quad 486, \quad 72,$$

with $\tau = 3,\ 5,\ 16,\ 125,\ 81,\ 12$ *respectively.*

Proof of Propositions 3.1, 3.2 and 3.4–3.8, and what was computed. Each statement is a finite evaluation of the enumerator of this section, carried out inside the proof assistant; Section 8 says exactly which evaluation belongs to which statement. Three of them are exhaustive over *all* matchings and not only the acyclic ones: the oracle agreement of Proposition 3.1 (every matching of four complexes, the largest being Δ^3 with 15 cells and 28 Hasse edges) and the two censuses of Proposition 3.5. In Proposition 3.8 the two sides are computed by algorithms sharing no code: the left by the matching search, the right by scanning all 2^m subsets of the edge set and keeping those with $n-1$ edges that connect the vertex set. □

4. The projective plane

Proof of Theorem 1.3, and what was computed. $\mathbb{RP}^2_6$ has 31 cells, so a matching has at most 15 pairs. Two exhaustive searches were run on its Hasse diagram, both testing acyclicity only at the leaves and pruning by acyclicity nowhere:

- all matchings with 15 pairs. There are 70 144 of them; the search tree has 1 983 447 nodes. *None* is acyclic.
- all matchings with 14 pairs. There are 1 698 480 of them; the search tree has 33 401 471 nodes. Exactly 428 400 are acyclic.

(The two matching totals and the node counts are instrumentation of an independent implementation of the same recursion and are reported for reproducibility; the machine-checked content is the two acyclic counts, 0 and 428 400.) The first search shows that no discrete Morse matching on $\mathbb{RP}^2_6$ has 15 pairs; the second exhibits 428 400 with 14. By (2) the number of critical cells is $31 - 2|W|$, hence odd, so the minimum is $31 - 2 \cdot 14 = 3$ and the matchings attaining it are exactly the 428 400 acyclic matchings of size 14. $\square$

Proof of Theorem 1.4, and what was computed. By Lemma 2.1 and (3) the Morse vector of a matching with three critical cells is one of $(1,1,1)$, $(2,1,0)$, $(0,1,2)$. Three profile-constrained searches, again testing acyclicity only at the leaves, give

$$\#\{(1,1,1)\} = 428\,400, \qquad \#\{(2,1,0)\} = 0, \qquad \#\{(0,1,2)\} = 0,$$

so every optimal matching has Morse vector $(1,1,1)$, and the first count agrees with the size-constrained count of Theorem 1.3 — two searches over different trees returning the same number. Since $\beta(\mathbb{RP}^2;\mathbb{F}_2) = (1,1,1)$, the optimal matchings are exactly the $\mathbb{F}_2$-perfect ones and $p_{\mathbb{F}_2}(\mathbb{RP}^2_6) = 428\,400$. Since $\beta(\mathbb{RP}^2;\mathbb{Q}) = (1,0,0)$, a rationally perfect matching would have a single critical cell, i.e. 15 pairs, and Theorem 1.3 rules that out; independently, a fourth profile-constrained search over the matchings with Morse vector $(1,0,0)$ returns 0. Hence $p_{\mathbb{Q}}(\mathbb{RP}^2_6) = 0$. $\square$

Remark 4.1. The two vanishing profiles of Theorem 1.4 also follow from Forman's weak Morse inequalities over $\mathbb{F}_2$: $c_i \geq \beta_i(\mathbb{RP}^2;\mathbb{F}_2) = 1$ for $i = 0,1,2$ excludes both $(2,1,0)$ and $(0,1,2)$. We compute them anyway, and record the agreement here rather than suppress it: it is a control on the profile-constrained search, whose other three outputs — 428 400, 0 for $(1,0,0)$, and the reproduction $[z_0 z_2]\,\mathrm{ME}_{\partial\Delta^3} = 256$ of Proposition 3.7 — are not independently predicted. The content of Theorem 1.4 that no inequality supplies is the value 428 400 of the perfect coefficient.

Remark 4.2. An explicit optimal matching is checked in the formal development as well; it is the control against the way a collection of counts can fail silently, by returning 0 everywhere. On $\mathbb{RP}^2_6$ the matching consists of the fourteen cover pairs

$$\begin{array}{llll}
\{1\}\text{–}\{0,1\}, & \{2\}\text{–}\{0,2\}, & \{3\}\text{–}\{0,3\}, & \{4\}\text{–}\{0,4\}, \\
\{5\}\text{–}\{0,5\}, & \{1,2\}\text{–}\{0,1,2\}, & \{1,3\}\text{–}\{0,1,3\}, & \{1,4\}\text{–}\{1,3,4\}, \\
\{2,4\}\text{–}\{0,2,4\}, & \{3,4\}\text{–}\{2,3,4\}, & \{1,5\}\text{–}\{1,2,5\}, & \{2,5\}\text{–}\{2,3,5\}, \\
\{3,5\}\text{–}\{0,3,5\}, & \{4,5\}\text{–}\{0,4,5\}, & &
\end{array}$$

and the checked assertions are that this list is a matching on $H(\mathbb{RP}^2_6)$, that it has fourteen pairs, that it is acyclic, that its Morse vector is $(1,1,1)$ and that its critical cells are the vertex $\{0\}$, the edge $\{2,3\}$ and the triangle $\{1,4,5\}$.

5. The torus and the dunce hat

The same instrument reaches two further complexes. Neither of the following is new; both are stated because they are the exact half of the torus and dunce-hat picture that is proved rather than measured, and because the counts of Remark 5.3 are meaningless without them.

Proposition 5.1. *Let T^2_7 be the 7-vertex triangulation of the torus given by* (4)*, with* 42 *cells. No matching of size* 21 *on $H(T^2_7)$ is acyclic and no matching of size* 20 *is acyclic; the two searches enumerated* 193 548 *and* 17 116 806 *matchings respectively, those totals being instrumentation and the two counts of acyclic matchings, both* 0*, being the machine-checked content. An explicit acyclic matching with* 19 *pairs, whose critical cells are the vertex $\{0\}$, the edges $\{2,3\}$ and $\{1,4\}$ and the triangle $\{3,4,5\}$, is exhibited and checked. Hence the minimum number of critical cells of a discrete Morse matching on T^2_7 is exactly* 4*, attained with Morse vector* $(1,2,1)$.

The value 4 is known: $(1,2,1)$ is the Betti vector of the torus over every field, so the weak Morse inequalities give ≥ 4, and a triangulated closed surface admits a matching realising its $\mathbb{F}_2$-Betti vector. What the proposition adds is that the bound is re-derived by exhaustion — [5, Tables 1–2] carry no torus row — and, with Remark 5.3, that T^2_7 is a complex where "optimal" and "perfect" agree over every field, in contrast with $\mathbb{RP}^2_6$.

Proposition 5.2. *Let D be Hachimori's 8-vertex triangulation of the dunce hat* [4]*, with facets*

013, 016, 017, 023, 024, 025, 045, 067, 124, 126, 127, 134, 237, 256, 345, 357, 567

on $\{0,\dots,7\}$. Then D has f-vector $(8,24,17)$, 49 cells, 99 Hasse edges and $\chi = 1$ — the row `dunce` of [5, Tables 1–2]*. No matching of size 24 on $H(D)$ is acyclic; the search enumerated all 15 529 992 of them, again with the total as instrumentation and the acyclic count 0 as the machine-checked content. An explicit acyclic matching with 23 pairs, with critical cells the vertex $\{0\}$, the edge $\{1,2\}$ and the triangle $\{5,6,7\}$, is exhibited and checked. Since $49-2|W|$ is odd, the minimum number of critical cells on D is exactly 3: the dunce hat is not collapsible, and its optimum value is $c = 3$.*

Non-collapsibility of the dunce hat is Zeeman's [8], and $c = 3$ is the value [5] obtain for `dunce` by branch and cut. Proposition 5.2 re-derives both from the definition of an acyclic matching, by exhaustion instead of by integer programming.

Remark 5.3 (measurements outside the formal development)*.* The *numbers* of optimal matchings on T^2_7 and on D were computed by enumerators written outside the proof assistant and are not part of the formal development. We report them here, labelled as measurements:

$$\#\{\text{optimal matchings of } T^2_7\} = 59\,523\,534, \qquad \#\{\text{optimal matchings of } D\} = 63\,617\,968,$$

every one of them with Morse vector $(1,2,1)$ and $(1,1,1)$ respectively — for the torus the other two available profiles are empty (218 169 966 matchings with profile $(0,2,2)$ and 16 846 872 with $(2,2,0)$ were enumerated and none is acyclic), and for the dunce hat the size-constrained and the $(1,1,1)$-constrained counts agree. For the torus the Morse inequalities already force the profile, $(1,2,1)$ being the Betti vector of T^2 over every field; for the dunce hat, which is contractible, the profile $(2,1,0)$ is excluded by no inequality and it is the agreement of the two counts that rules it out. Four enumerations agree on 59 523 534 and three on 63 617 968: a size-constrained leaf-only search, a profile-constrained leaf-only search, an acyclicity-pruned search (torus only) and a level-factorised algorithm that decomposes an acyclic matching on a 2-complex into its vertex–edge and edge–triangle levels and recombines them over edge sets. The first three are different searches within one program and share its construction of the Hasse diagram and its acyclicity test, differing in what they enumerate and where acyclicity is tested; the last shares no code with them. Its correctness rests on the observation that every directed cycle of $H_W(K)$ alternates up- and down-steps and therefore lives in two consecutive dimensions, which we have not formalised. That is precisely why these two numbers are measurements here and the 428 400 of Theorem 1.3 is not.

6. Costs, and what was not run

This section reports measurements; nothing else in the paper depends on it.

Inside the proof assistant, with the enumerator compiled (Section 8), the whole development costs about 395 CPU-seconds on one core of a shared 20-core machine: the controls of Section 3 2.95 s, the Δ^4 reproductions 28.25 s, Theorems 1.3 and 1.4 together 36.69 s, Proposition 5.1 169.15 s and Proposition 5.2 158.13 s, with a peak resident set of about 1.5 GB. The same searches in a C implementation run 24 to 26 times faster; for instance the size-14 search on $\mathbb{RP}^2_6$ takes 0.50 CPU-seconds there.

Two consequences worth recording. First, on the reproduction of $f(4)$: the top entry alone — which is the published number 380 125 — costs 0.36 CPU-seconds of C on one core, and 28 CPU-seconds inside the proof assistant. This is a remark about implementations and not about the mathematics: the cluster of [7] was needed for the *full* 15-entry f-vector of $M(\Delta^4)$, whose largest entry is 876 110 235, and the reproduction repository accompanying that paper itself recommends porting its reference implementation to a compiled language.

Second, on what was not run. Promoting the two measurements of Remark 5.3 into the formal development by the same route is a matter of cost and not of method: the profile-constrained search for the torus is 67.11 CPU-seconds in C, hence about 1 750 CPU-seconds inside the proof assistant, and the dunce hat about 3 300 CPU-seconds — both within reach, both past the wall-clock budget we allow a single file, and both splittable into seven or eight independent sub-counts by the branch taken at the first vertex. We record the price rather than pay it here. The next value $f(5)$ of the simplex sequence is out of reach by this method: [7, §5.2] calls even the f-vector of $M(\Delta^5)$ computationally infeasible, and although the top entry alone is far cheaper than the whole vector — that is the $n = 4$ lesson above — extrapolating the searches timed here still puts $f(5)$ well past the compute budget we allow a single family. So, for now, is the count for Bing's house, one of the five complexes of Hachimori's library that the branch and cut of [5] could not solve in reasonable time.

7. A note on the source

The paper that poses the counting question in the form we take up is unusually explicit about its own provenance. [7] closes with a section headed "AI Acknowledgments" which records the use of Anthropic's Claude Opus 4.6, states that it "served as a dialogue partner during the development of proofs, flagged gaps and errors in arguments, helped structure and edit the exposition, and was instrumental in writing and validating the enumeration scripts, Smith normal form computations, and the paper's verification software", and adds that "the author retains full responsibility for the correctness of all mathematical content; every proof and computation has been independently checked", citing a public repository that "independently reproduces every numerical claim". We record this as a matter of fact and take no position on it; the numbers of [7] that we could reach are reproduced above and all of them agree.

The difference the present paper adds is not in how the numbers were found. That repository is a Python test suite, reporting 51 independent checks over Δ^n, $\partial\Delta^n$ and their pure and generalised variants; its files are Python and it contains no proof-assistant code, so what is attached to its numbers is a passing run. Here the reproductions are theorems that a proof assistant's kernel accepts, together with negative controls that fail (Section 3), and the same development that reproduces 380 125 produces 428 400.

8. Verification

Machine-checked in Lean 4 [6] are Theorems 1.3 and 1.4 and Propositions 3.1, 3.2, 3.4–3.8, 5.1 and 5.2, in the form stated, together with the explicit optimal matchings of Remark 4.2 and Propositions 5.1, 5.2: twenty-two theorems in all. Of Proposition 3.2 it is the combinatorial half that is checked — closed surface, f-vector, Euler characteristic, orientability, and the non-vacuity of the test — while the step from those invariants to the names $\mathbb{RP}^2$, T^2 and S^2, namely the classification of closed surfaces, is not. Not formalised either are Lemma 2.1 and (2), which are the elementary counting proved above; the two classical facts quoted in Section 2; the census of Remark 3.3; the measurements of Remark 5.3 and Section 6; and Remark 4.1, which is commentary.

The development imports no mathematical library, Mathlib included. A complex is a record of flat arrays of natural numbers — cell dimensions, vertex bitmasks, covering and coboundary lists, and the successor lists of the modified Hasse diagram — built from the facet list by a single function, so the object the kernel checks is the Hasse diagram itself and every value in sight is a natural number or an array of them. The finite evaluations are carried out by Lean's compiled evaluator (`native_decide`); each of the twenty-two theorems therefore depends on `propext`, `Quot.sound` and one evaluation axiom of its own, and that evaluator is the only trusted step. There is no `sorry` anywhere in the development, it declares no axiom of its own, and it uses no form of the axiom of choice. Because the searches are large, one invocation of the Lean frontend on the module holding the enumerator emits both the compiled artifact the kernel checks and the C source of that same module; the C is compiled and linked into a shared library, and the remaining modules are checked against that library, which keeps the evaluations compiled rather than interpreted (a factor of about 24 on this development). Since the two outputs come from one invocation on one source file, the library the evaluations run and the module the kernel checks are the same module. Independently written C programs reproduce every count reported above, and for the counts of Remark 5.3, which the proof assistant does not certify, four respectively three enumerations agree, in the sense made precise there. The repository is private; repository reference to be added at submission.

References

[1] B. Benedetti and F. H. Lutz, *Random discrete Morse theory and a new library of triangulations*, Experimental Mathematics **23** (2014), no. 1, 66–94; `doi:10.1080/10586458.2013.865281`; arXiv:1303.6422.

[2] M. K. Chari and M. Joswig, *Complexes of discrete Morse functions*, Discrete Mathematics **302** (2005), no. 1–3, 39–51; `doi:10.1016/j.disc.2004.07.027`. The source of $f(1) = 2$, $f(2) = 9$, $f(3) = 256$ and of the complex $M(K)$ used above.

[3] R. Forman, *Morse theory for cell complexes*, Advances in Mathematics **134** (1998), no. 1, 90–145; `doi:10.1006/aima.1997.1650`.

[4] M. Hachimori, *Simplicial Complex Library*, entry "Dunce hat" (8 vertices, 17 facets, $f = (1, 8, 24, 17)$), with facet list in the linked file `dunce_hat.dat`, `http://infoshako.sk.tsukuba.ac.jp/~hachi/math/library/dunce_hat_eng.html`; retrieved 3 September 2026. The facet list of Proposition 5.2 is that file's, re-indexed from $1, \ldots, 8$ to $0, \ldots, 7$.

[5] M. Joswig and M. E. Pfetsch, *Computing optimal Morse matchings*, SIAM Journal on Discrete Mathematics **20** (2006), no. 1, 11–25; `doi:10.1137/S0895480104445885`; arXiv:math/0408331. The tables cited above are Tables 1 and 2 of the arXiv version, the version we read; the rows used are `projective` ($n = 31$, $m = 60$, $\beta = 3$, $c = 3$) and `dunce` ($n = 49$, $m = 99$, $\beta = 1$, $c = 3$), where in their notation n is the number of faces, m the number of arcs of the Hasse diagram, β the lower bound from the Betti inequalities and c the number of critical faces in the optimal solution. Their test set is Hachimori's [4].

[6] L. de Moura and S. Ullrich, *The Lean 4 theorem prover and programming language*, in: A. Platzer and G. Sutcliffe (eds.), Automated Deduction — CADE 28, Lecture Notes in Computer Science **12699**, Springer, Cham, 2021, pp. 625–635; `doi:10.1007/978-3-030-79876-5_37`.

[7] N. A. Scoville, *The complex of discrete Morse matchings of the n-simplex: homotopy types and structural results*, arXiv:2604.18172v1 [math.AT], submitted 20 April 2026; v1 is the only version at the time of writing, and every pointer to it, here and in the text above, is to that version. Section 2 fixes the definitions used in Section 1; the sentence quoted there opens Section 3; Table 1 gives the two f-vectors of Proposition 3.5; Section 5.2 carries the f-vector of $M(\Delta^4)$ and the remark on $n = 5$; Proposition 9 is the top-dimensional-facet bijection used in Proposition 3.6; the closing section is headed "AI Acknowledgments".

[8] E. C. Zeeman, *On the dunce hat*, Topology **2** (1963), 341–358; `doi:10.1016/0040-9383(63)90014-4`. Crossref and zbMATH both date this to 1963; it is often cited as 1964.

[9] C. Zheng, *On the Morse ensemble polynomial of simplicial complexes*, arXiv:2605.24689v3 [math.CO], revised 28 June 2026. Every pointer to it, here and in the text above, is to v3, whose numbering need not agree with v1 or v2. Definition 1.1 is the Morse ensemble polynomial; Section 2 defines perfect matchings and Section 5.5 the perfect coefficients $p_k(K)$; Theorem 4.9 is $p(G) = n\,\tau(G)$ for connected graphs and Example 4.6 the cospectral pair G_1/G_2, both reproduced in Propositions 3.7 and 3.8; the remark on matroid independence complexes in Section 6 states $\mathrm{ME}_{\mathrm{Ind}(U_{2,4})}$ and $[z_0 z_2]\,\mathrm{ME}_{\partial\Delta^3} = 256$; the open problems are the enumerated list of Section 7, and "Perfect Morse counts in higher dimensions" is its second item.