



HAL
open science

Slit-slide-sew bijections for constellations and quasiconstellations

Jérémie Bettinelli, Dimitri Korkotashvili

► **To cite this version:**

Jérémie Bettinelli, Dimitri Korkotashvili. Slit-slide-sew bijections for constellations and quasiconstellations. 2023. hal-04299507

HAL Id: hal-04299507

<https://hal.science/hal-04299507>

Preprint submitted on 11 Mar 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Slit-slide-sew bijections for constellations and quasiconstellations

J eremie Bettinelli*

Dimitri Korkotashvili[†]

March 11, 2024

Abstract

We extend so-called slit-slide-sew bijections to constellations and quasiconstellations. We present an involution on the set of hypermaps given with an orientation, one distinguished corner, and one distinguished edge leading away from the corner while oriented in the given orientation. This involution reverts the orientation, exchanges the distinguished corner with the distinguished edge in some sense, slightly modifying the degrees of the incident faces in passing, while keeping all the other faces intact.

The construction consists in building a canonical path from the distinguished elements, slitting the map along it, and sewing back after sliding by one unit along the path. The involution specializes into a bijection interpreting combinatorial identities linking the numbers of constellations or quasiconstellations with a given face degree distribution, where the degree distributions differ by one $+1$ and one -1 . In particular, this allows to recover the counting formula for constellations or quasiconstellations with a given face degree distribution.

Our bijections furthermore provide an algorithm for sampling a hypermap uniformly distributed among constellations or quasiconstellations with prescribed face degrees.

1 Introduction

In the present work, we pursue the investigation of so-called *slit-slide-sew* bijections, introduced in [Bet14] on forests and plane quadrangulations, and then further developed in [Bet19, Bet20] on plane bipartite and quasibipartite maps. Here, we focus on a generalization of the latter, called *constellations* and *quasiconstellations*.

Hypermaps. Recall that a *plane map* is an embedding of a finite connected graph (possibly with multiple edges and loops) into the sphere, considered up to orientation-preserving homeomorphisms. Now fix an integer $p \geq 2$. A (plane) *p-hypermap* is a plane map whose faces are shaded either *dark* or *light* in such a way that

- ◊ adjacent faces do not have the same shade (one is dark, the other light);

*cnrs & Laboratoire d'Informatique de l' cole polytechnique; jeremie.bettinelli@normalesup.org; www.normalesup.org/~bettinel. Partially supported by Grant ANR-20-CE48-0018 *3DMaps*.

[†]Laboratoire d'Informatique de l' cole polytechnique; dimitri.korkotashvili@polytechnique.edu.

◇ each dark face has degree p .

These actually generalize maps, which correspond to 2-hypermaps. In the terminology of hypermaps, light faces generalize faces and might be called *hyperfaces*, whereas dark faces generalize edges and are called *hyperedges*. We do not use this terminology here.

A (plane) *p*-constellation is a p -hypermap such that the degrees of its light faces are all multiples of p . In a *p*-hypermap, a light face whose degree is not a multiple of p will be called a *flawed face*. A *p*-constellation is thus a p -hypermap without flawed faces. A *quasi-p-constellation* is a p -hypermap with exactly two flawed faces. Note that, in a p -hypermap, the sum of the degrees of the light faces is necessarily a multiple of p , since it is equal to the sum of the degrees of the dark faces, which are all p . As a result, a p -hypermap cannot have a single flawed face and, in a quasi- p -constellation, the two flawed faces have, modulo p , degrees $+k$ and $-k$ for some $0 < k < p$.

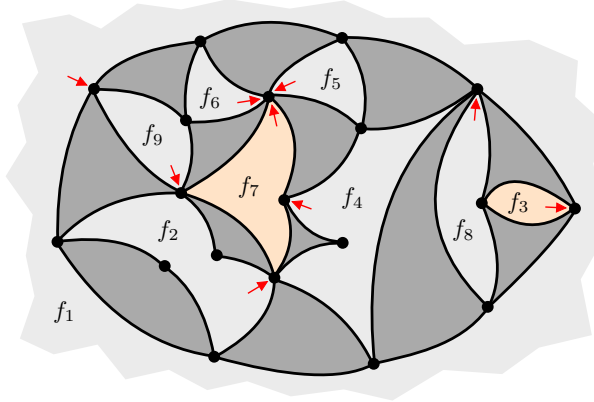


Figure 1: A quasi-3-constellation of type $(9, 6, 2, 6, 3, 3, 4, 3, 3)$. The two flawed faces are f_3 and f_7 ; throughout the paper, we highlight flawed faces in orange. Every light face has a marked corner, always represented by a red arrowhead.

In the particular case $p = 2$, squeezing each dark face (which has degree 2 by definition) into a single edge yields a simple one-to-one correspondence between 2-constellations and bipartite maps, as well as between quasi-2-constellations and quasibipartite maps, as shown in Figure 2.

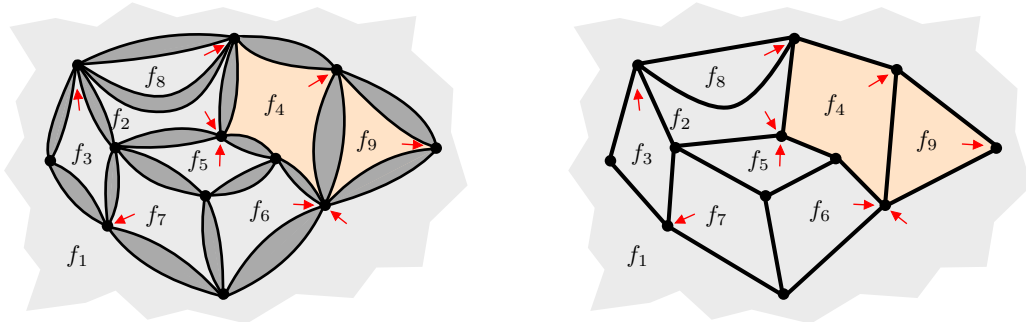


Figure 2: A quasi-2-constellation and the corresponding quasibipartite map, obtained by squeezing all the dark faces.

Another particular class of interest enters our framework: the class of Eulerian triangulations, that is, 3-constellations whose light faces are all of degree 3.

Enumeration. For an r -tuple $\mathbf{a} = (a_1, \dots, a_r)$ of positive integers, let us denote by $C(\mathbf{a})$ the number of p -hypermaps with exactly r light faces, numbered f_1, \dots, f_r and of respective degrees a_1, \dots, a_r , each bearing a marked corner¹. The r -tuple \mathbf{a} will be called the *type* of such p -hypermaps. See Figure 1. By elementary considerations and Euler's characteristic formula, the integers

$$E(\mathbf{a}) := \sum_{i=1}^r a_i, \quad D(\mathbf{a}) := \frac{E(\mathbf{a})}{p}, \quad \text{and} \quad V(\mathbf{a}) := E(\mathbf{a}) - D(\mathbf{a}) - r + 2$$

are respectively the numbers of edges, dark faces, and vertices of p -hypermaps of type \mathbf{a} . Generalizing Tutte's so-called *formula of slicings* [Tut62], it has been computed that, when at most two a_i 's are not in $p\mathbb{N}$, that is, for p -constellations [BMS00] or quasi- p -constellations [CF14], it holds that

$$C(\mathbf{a}) = c_{\mathbf{a}} \frac{(E(\mathbf{a}) - D(\mathbf{a}) - 1)!}{V(\mathbf{a})!} \prod_{i=1}^r \alpha(a_i), \quad \text{where} \quad \alpha(x) := \frac{x!}{[x/p]! (x - [x/p])!} \quad (1)$$

$$\text{and} \quad c_{\mathbf{a}} = \begin{cases} 1 & \text{if } p \text{ divides every } a_i \\ p-1 & \text{otherwise} \end{cases}.$$

Combinatorial identities. In the present work, we give a bijective interpretation for the following combinatorial identity, which transfers one degree from one face to another.

Proposition 1 (Transferring one degree from f_1 to f_2). *Let $\mathbf{a} = (a_1, \dots, a_r)$ be an r -tuple of positive integers such that $a_1 \geq 2$, and with coordinates equal modulo p to*

- (i) *either $(k, -k, 0, \dots, 0)$ for some $k \in \{0, \dots, p-1\}$,*
- (ii) *or $(1, 0, \dots, 0, -1, 0, \dots, 0)$, with the -1 in any position from 3 to r .*

Let also $\tilde{\mathbf{a}} = (\tilde{a}_1, \dots, \tilde{a}_r) := (a_1 - 1, a_2 + 1, a_3, \dots, a_r)$. Then the following identity holds:

$$(a_1 - \lceil a_1/p \rceil) (a_2 + 1) C(\mathbf{a}) = (\tilde{a}_1 + 1) (\tilde{a}_2 - \lceil \tilde{a}_2/p \rceil) C(\tilde{\mathbf{a}}). \quad (2)$$

To obtain (2) from (1), one might first observe that, for any $x \in \mathbb{N}$,

$$\frac{\alpha(x)}{\alpha(x-1)} = d_x \frac{x}{x - \lceil x/p \rceil} \quad \text{where} \quad d_x = \begin{cases} p-1 & \text{if } p \mid x \\ 1 & \text{if } p \nmid x \end{cases},$$

and then that, in both cases (i) and (ii), $c_{\mathbf{a}} d_{a_1} = c_{\tilde{\mathbf{a}}} d_{\tilde{a}_2} = p-1$.

Remark 1. In fact, Proposition 1 also holds when the coordinates of \mathbf{a} are equal modulo p to

- (iii) $(0, -1, 0, \dots, 0, 1, 0, \dots, 0)$, with the $+1$ in any position from 3 to r ,

¹Recall that a *corner* is an angular sector delimited by two consecutive edges around a vertex.

observing that, in this case, $c_a d_{a_1} = c_{\tilde{a}} d_{\tilde{a}_2} = (p-1)^2$. For some technical reason, however, this particular case falls out of the reach of the present paper.

We furthermore treat the case of a degree 1-face, which may easily be obtained as above.

Proposition 2 (Transferring the degree of a degree 1-face f_1 to f_2). *Let $\mathbf{a} = (1, a_2, \dots, a_r)$ and $\tilde{\mathbf{a}} = (\tilde{a}_2, \dots, \tilde{a}_r) := (a_2 + 1, a_3, \dots, a_r)$ be respectively an r -tuple and an $r - 1$ -tuple of positive integers, both having at most two coordinates not lying in $p\mathbb{N}$. Then the following identity holds:*

$$(a_2 + 1) C(\mathbf{a}) = V(\tilde{\mathbf{a}}) (\tilde{a}_2 - \lceil \tilde{a}_2/p \rceil) C(\tilde{\mathbf{a}}). \quad (3)$$

It is easy to see that the number of p -constellations with exactly one light face of degree pn is equal to the number of p -ary trees with n nodes: see Figure 3. It thus holds (see for instance [GJ83, Section 2.7.7.]) that

$$C((pn)) = \frac{(pn)!}{n! ((p-1)n + 1)!}.$$

Using this formula as initial condition, Propositions 1 and 2 provide yet another proof of (1).

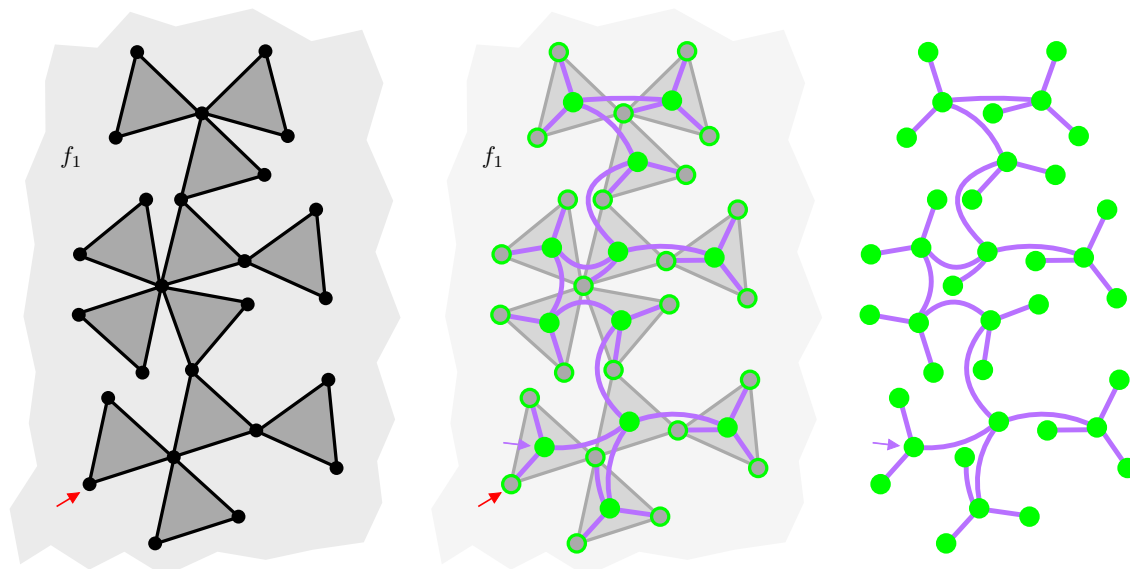


Figure 3: Bijection between p -constellations of type (pn) and p -ary trees with n nodes. Each dark face corresponds to a p -ary node and each vertex to a leaf. We follow the contour of the light face, starting from its marked corner, in a given orientation, say with the light face on the left. Each time we encounter an unvisited dark face, say D , we consider the 3 incident vertices, say v_1, v_2, v_3 . For $1 \leq i \leq 3$, we link the node corresponding to D to the node corresponding to the left-most unvisited dark face incident to v_i or, by default, to the leaf corresponding to v_i . We root the tree at the corner following the marked corner of f_1 .

Methodology. In order to bijectively interpret the combinatorial identities (2) and (3), the idea is to distinguish elements, such as edges, vertices, faces, corners, etc., in such a way that each side of

an equation of interest counts maps given with such distinguished elements. For instance, in the left-hand side of (2), the term $a_2 + 1$ counts the number of corners² in the face f_2 of a p -hypermap of type \mathbf{a} . As a result, $(a_2 + 1) C(\mathbf{a})$ is the number of p -hypermaps of type \mathbf{a} with a distinguished corner in the face f_2 . The other prefactor in the left-hand side of (2) is more delicate to interpret; this will be done in due course. Remark that we will always use the word “distinguished” to designate these extra elements, keeping the word “marked” only for the marked corners, which we see as inherent to the hypermaps into consideration.

Once both sides of the considered equation are properly interpreted as cardinalities of sets of maps with distinguished elements, we bijectively go from one set to the other as follows. Using the distinguished elements, we construct a directed path in the map, called *sliding path*. We then slit the map along this sliding path and sew back together the sides of the slit after sliding by one unit, in the sense that the left side of the i -th edge is sewn back on the right side of the $i \pm 1$ -th edge (the ± 1 being the same for all edges and determined by some rule). This mildly modifies the map along the path but does not affect its faces, except the two that are around the extremities of the sliding path. In the process, new distinguished elements naturally appear in the resulting map; these allow us to recover the sliding path in order to slide back.

Organization of the paper. The remainder of the document is structured in the following manner. We start by giving in Section 2 the definitions and conventions we use, as well as a combinatorial interpretation of the prefactor $(a - \lceil a/p \rceil)$ appearing in the identities (2) and (3). We then present in Section 3 our bijective interpretation of these identities through a more general involution on the set of maps given with an orientation, a distinguished corner, and a distinguished edge satisfying an extra constraint. Finally, we explain in Section 4 how to sample a uniform p -hypermap of a given type using our bijections.

2 Preliminaries

Distinguishing a corner. Following previous works on slit-slide-sew bijections, we use the convention, depicted in Figure 4, that the marked corner of a face creates two possible corners to distinguish. One might think of the marked corner as a dangling half-edge, with one corner on each side. As a result, a face of degree a bearing its marked corner has $a + 1$ possible corners to distinguish.

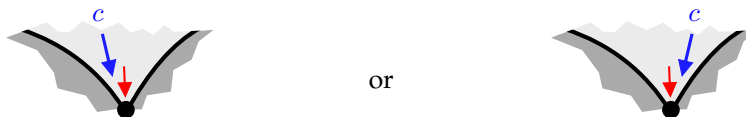


Figure 4: Distinguishing a corner around the marked corner.

Edge orientation. As is customary, we will orient the edges of the hypermaps we consider, in such a way that light faces always lie to the same side of the oriented edges (and thus dark faces

²The $+1$ comes from the presence of the marked corner in the face. See the first paragraph of Section 2 for the convention we use here.

always lie to the same other side). These orientations will be called the *light-left orientation* when the light faces³ all lie to the left, and the *light-right orientation* when the light faces all lie to the right. In other words, in the light-right orientation, the edges are oriented clockwise around light faces and counterclockwise around dark faces. See Figure 5. We will need to use both orientations in the present paper. We will always clearly mention which orientation we use whenever it matters. Without specific mention, both orientations can be used. **Once one of the two possible orientations is fixed**, we will use the following conventions.

Given an edge e , we will respectively denote by e^- and e^+ the origin and end of the edge e , oriented as convened. The corner *preceding* e is defined as the corner c_e delimited by e and the edge that precedes e in the contour of the incident light face, in the convened orientation. Similarly, we denote by c^+ the vertex incident to a corner c .

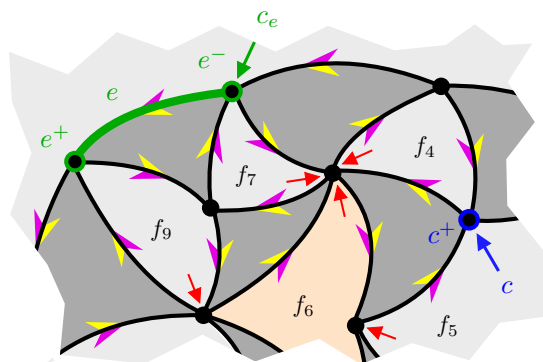


Figure 5: Edge orientation and related definitions. Here, the light-right orientation is depicted.

Paths. A *path* from a vertex v to a vertex v' is a finite sequence $\wp = (e_1, e_2, \dots, e_k)$ of edges such that $e_1^- = v$, for $1 \leq i \leq k-1$, $e_i^+ = e_{i+1}^-$, and $e_k^+ = v'$. Its *length* is the integer k , which we denote by $[\wp] := k$. A path is called *simple* if the vertices it visits are all distinct.

Beware that a path is only made of edges oriented in the convened orientation. In other words, edges cannot be used “backward.” In particular, this means that all the faces lying to the left of a path are of the same shade (either all light or all dark), whereas all the faces lying to its right are of the other shade. The side of the path where the faces are all light will be called its *light side*, whereas the other side will be called its *dark side*.

Directed metric and geodesics. We will use the directed metric associated with the convened orientation: given two vertices v, v' in a p -hypermap, we denote by $\vec{d}(v, v')$ the smallest k for which there exists a path from v to v' of length k . (We put an arrow on top in the notation to keep in mind that this is only a directed metric.) A *geodesic* from v to v' is such a path.

There are generally several geodesics from a given vertex v to a target vertex v' . Among all of these, one will be of particular interest in this work: the *lightest geodesic*, constructed as follows. It is only well defined from a starting edge or corner e_0 such that $e_0^+ = v$. (The starting element e_0

³Recall that the light faces are the main objects of focus.

does not belong to the path.) Then, provided e_0, e_1, \dots, e_j have already been constructed and the path is not complete (that is, $e_j^+ \neq v'$), we set the subsequent edge e_{j+1} as the one, among the edges e such that $e^- = e_j^+$ and $\vec{d}(e^+, v') = \vec{d}(e_j^+, v') - 1$, that comes first while turning around e_j^+ in the direction incoming edge, light face. See Figure 6. In other words, the lightest geodesic is the leftmost geodesic if the convened orientation is the light-left orientation and the rightmost geodesic if the convened orientation is the light-right orientation.

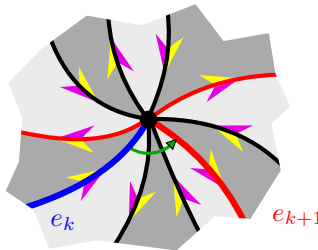


Figure 6: Definition of the lightest geodesic. The three red edges are the ones going closer to the target vertex v' . Turning around e_k^+ in the direction of the light face, we take the first red edge (drawn with a thicker line).

Edge types. Given a fixed vertex v in a p -hypermap, we may differentiate three types of edges: an edge e is said to be

- ◇ *leaving* v if $\vec{d}(v, e^+) = \vec{d}(v, e^-) + 1$;
- ◇ *approaching* v if $\vec{d}(v, e^+) = \vec{d}(v, e^-) + 1 - p$;
- ◇ *irregular with respect to* v if $\vec{d}(v, e^+) - \vec{d}(v, e^-) \not\equiv 1 \pmod{p}$.

Observe that $1 - p \leq \vec{d}(v, e^+) - \vec{d}(v, e^-) \leq 1$ since there is always a path of length 1, namely the path consisting of the single edge e , as well as a path from e^+ to e^- of length $p - 1$, made of all the other edges incident to the dark face incident to e . As a result, if e is irregular with respect to v , then it holds that $\vec{d}(v, e^+) - \vec{d}(v, e^-) \in \{2 - p, 3 - p, \dots, 0\}$. See Figure 7.

The following proposition gives the number of each type among edges incident to a given face in a p -constellation or a given flawed face in a quasi- p -constellation; this provides an interpretation to the prefactor $(a - \lceil a/p \rceil)$ appearing in (2) and (3).

Proposition 3. *We consider a vertex v and a light face f of degree a in a p -hypermap.*

- (1) *If the p -hypermap is a p -constellation then, among the a edges incident to f ,*
 - ◇ $a - a/p$ are leaving v ;
 - ◇ a/p are approaching v ;
 - ◇ none are irregular with respect to v .
- (2) *If the p -hypermap is a quasi- p -constellation and f a flawed face then, among the a edges incident to f ,*

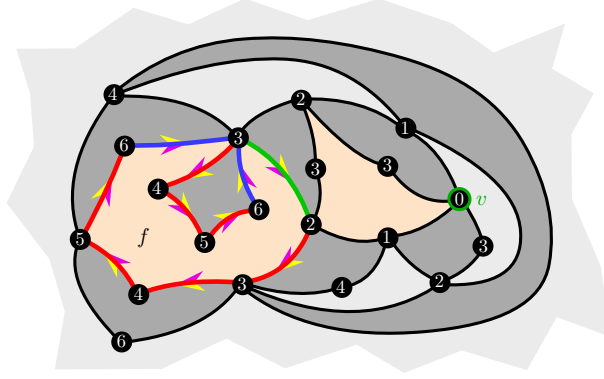


Figure 7: The different types of edges incident to a flawed face in a quasi-4-constellation. We use the light-right orientation; the distances to v are written in the vertices. Around f , the $(10 - \lceil 10/4 \rceil) = 7$ red edges are leaving v ; the $\lfloor 10/4 \rfloor = 2$ blue edges are approaching v ; the green edge is irregular with respect to v .

- ◇ $(a - \lceil a/p \rceil)$ are leaving v ;
- ◇ $\lfloor a/p \rfloor$ are approaching v ;
- ◇ one is irregular with respect to v .

Proof. First of all, observe that, for either (1) or (2), the third subitem easily implies the first and second subitems. Indeed, by definition and from the observation above, following the contour of f , the (oriented) distance variations to v are

- ◇ $+1$ for leaving edges;
- ◇ $+1 - p$ for approaching edges;
- ◇ $-k$ for some $k \in \{0, 1, \dots, p-2\}$ for the irregular edge, in the quasi- p -constellation case.

These differences must sum up to 0, which gives the desired result (and shows in passing that, in fact, $k = a - p\lfloor a/p \rfloor - 1$).

Now, note that, in the quasi- p -constellation case, there must be at least one irregular edge since, for all other edges, the distance variations to v are equal to $1 \pmod p$ and sum up to 0. The third subitem is thus obtained from the case $k = 0$ for (1) and $k = 2$ for (2) of the following more general claim.

Claim. We consider a vertex v and a light face f in a p -hypermap having exactly k flawed faces. Then, the number of edges incident to f that are irregular with respect to v is

- ◇ at most $k - 1$ if f is a flawed face;
- ◇ at most k otherwise.

To show this, let us denote by e_1, \dots, e_m the edges incident to f that are irregular with respect to v . Then, for each $1 \leq i \leq m$, let us consider two geodesics γ_i^- and γ_i^+ from v towards e_i^- and e_i^+ . Note that, since e_i is not leaving v , the geodesics γ_i^- and γ_i^+ have to split before reaching e_i^- .

Up to changing the geodesics, one may furthermore assume that any two geodesics never meet again after the time they split. These $2m$ geodesics make up a directed tree whose leaves are the extremities of e_1, \dots, e_m . See Figure 8.

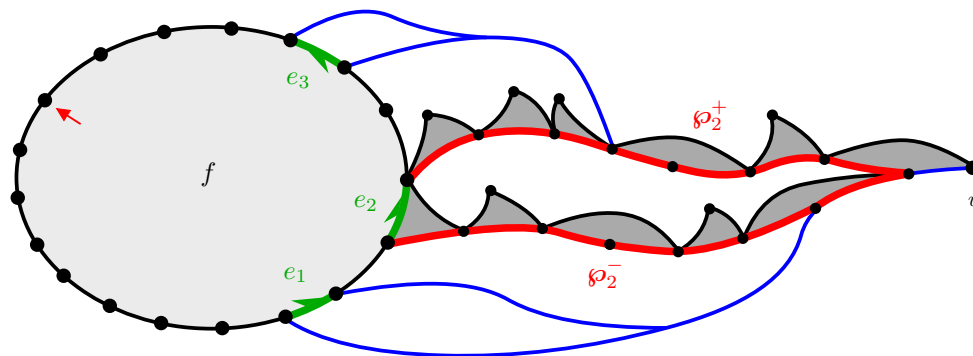


Figure 8: The face f and the irregular edges e_1, e_2, e_3 . The geodesics are in blue, except the parts \wp_2^- and \wp_2^+ , which are highlighted in red. Only the dark faces that are to the dark side of \wp_2^- and of \wp_2^+ are represented.

Now, for each i , let us consider the parts \wp_i^- and \wp_i^+ of γ_i^- and γ_i^+ after the vertex where they split up. The concatenation of \wp_i^- , e_i , and then the reverse of \wp_i^+ provides a cycle of edges (not all oriented in the same direction) in the p -hypermap. By the Jordan curve theorem, this cycle separates the sphere into two connected components: let us denote by \mathcal{C}_i the one that does not contain the face f . Note that, by construction, $\mathcal{C}_1, \dots, \mathcal{C}_m, f$, are pairwise disjoint. Consequently, the proof of the claim will be complete if we show that each \mathcal{C}_i contains at least one flawed face, since this yields that $m + \mathbb{1}_{\{f \text{ is flawed}\}} \leq k$.

In order to see this, let us consider the p -hypermap constructed as shown in Figure 9. We start from \mathcal{C}_i and we add one dark face to the dark side of each of the $[\wp_i^+]$ edges on the part of the boundary of \mathcal{C}_i that corresponds to \wp_i^+ . Then we add a light face (of degree $[\wp_i^-] + 1 + (p-1)[\wp_i^+]$) on the boundary of the resulting object. (Recall that \wp_i^+ is taken backward in the cycle and thus has its incident dark faces out of \mathcal{C}_i .) The marked corner of the added light face is set arbitrarily.

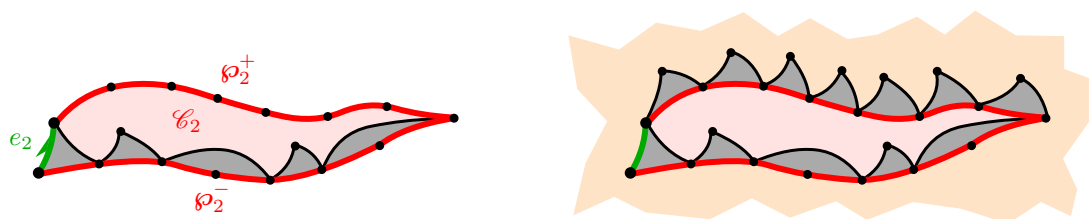


Figure 9: The connected component \mathcal{C}_2 and the corresponding p -hypermap (the irrelevant marked corner of the added light face is not depicted). The fact that the unbounded face is flawed implies the existence of a flawed face in the red area.

Finally, the added light face is necessarily a flawed face since, modulo p , its degree is equal to $[\wp_i^-] + 1 + (p-1)[\wp_i^+] \equiv \bar{d}(v, e_i^-) + 1 - \bar{d}(v, e_i^+) \not\equiv 0$, by definition of e_i being irregular. As a result,

there should be another flawed face in this p -hypermap, which is necessarily inside \mathcal{C}_i . The claim follows. \square

3 Bijjective interpretation

3.1 Slit slide sew

Let us first describe the operation at the heart of our construction. See Figure 10. Assume that, on some p -hypermap m , we have a simple path $\wp = (e_1, e_2, \dots, e_k)$ linking some corner c in some light face f to a different corner c' in some light face f' (which may possibly be equal to f), that is, such that $e_1^- = c^+$ and $e_k^+ = c'^+$. We may then follow \wp , **entering from the corner c and exiting through the corner c'** . This creates a simple path on the sphere, starting inside the face f and finishing inside f' . We may slit the sphere along this path, thus doubling the sides of the path. In the hypermap m , this doubles the path \wp , making up two copies, one incident to light faces and the slit, and one incident to dark faces and the slit. We denote by $\ell = (\ell_1, \dots, \ell_k)$ the former and by $d = (d_1, \dots, d_k)$ the latter.

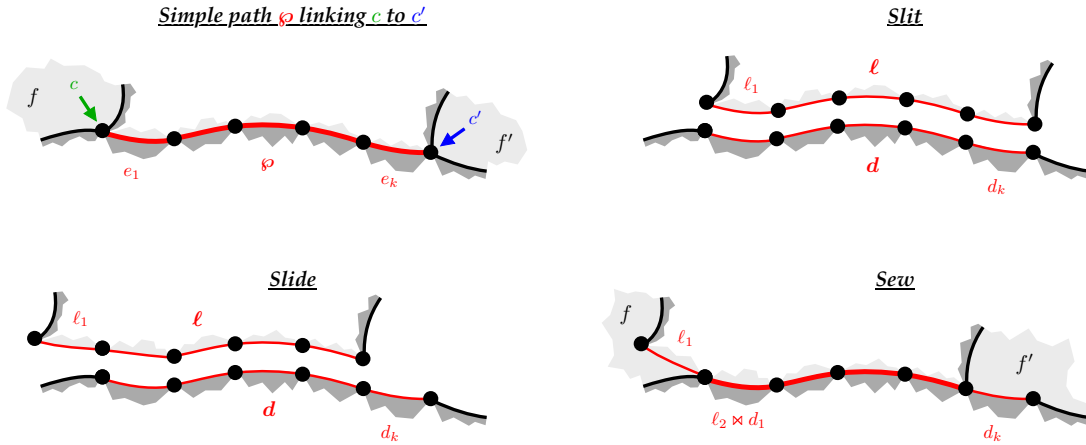


Figure 10: The slit-slide-sew operation on a p -hypermap. Note that, here, ℓ_1 is not a dangling edge since c is not the corner preceeding e_1 .

Note that the data of \wp is not sufficient to properly define this operation; one needs to know from which corner to enter \wp in order to decide if an edge incident to e_1^- becomes incident whether to ℓ_1^- or to d_1^- . Similarly, one needs to know through which corner to exit \wp .

We then sew back ℓ onto d but only after sliding by one unit, in the sense that we match ℓ_{i+1} with d_i , for every $1 \leq i \leq k - 1$. For further reference, we denote by $\ell_{i+1} \bowtie d_i$ the resulting edge. Observe that, except from f and f' , the faces are not altered by the process. Observe also that ℓ_1 and d_k are not matched with anything:

- \diamond d_k is still incident to the original dark face and is now also incident to f' ;
- \diamond ℓ_1 is still incident to the original light face and is now also incident to f .

Consequently, the result is no longer a p -hypermap since ℓ_1 is incident to light faces from both sides. However, in the case where ℓ_1 is actually a dangling edge (an edge with one extremity of degree 1), removing it provides a p -hypermap. This happens if and only if c is the corner preceding e_1 ; this will always be the case in the present work.

3.2 Face of degree two or more

We now present the bijective interpretation for the identity (2) of Proposition 1.

Involution. We start by defining a mapping Φ on the set \mathcal{H} of quadruples $(\varphi, \mathfrak{m}, c, e)$, where

- ◇ φ is an orientation (either light-left or light-right);
- ◇ \mathfrak{m} is a p -hypermap;
- ◇ c is a distinguished corner of some light face;
- ◇ e is a distinguished edge leaving c^+ in the orientation φ .

We break down the process into the following steps. See Figure 11.

1. Reorientation

From this point on, we convene to use the reverse orientation, which we denote by $\tilde{\varphi}$.

2. Sliding path

We consider the corner c_e preceding e and the lightest geodesic γ from c_e to c^+ .

3. Slitting, sliding, sewing

We slit, slide, sew along γ from c_e to c as described in the previous section: along γ , the light side of an edge is now matched with the dark side of the previous edge.

4. Output

The unmatched light side of the first edge of γ yields a dangling edge; we remove it and denote the resulting corner by \tilde{c} . We denote the edge corresponding to the unmatched dark side of the final edge of γ by \tilde{e} . We let $\tilde{\mathfrak{m}}$ be the resulting map. Finally, the output of the construction is the quadruple $\Phi(\varphi, \mathfrak{m}, c, e) := (\tilde{\varphi}, \tilde{\mathfrak{m}}, \tilde{c}, \tilde{e})$.

Theorem 4. *The mapping $\Phi: \mathcal{H} \rightarrow \mathcal{H}$ is an involution.*

Proof. We consider an element $(\varphi, \mathfrak{m}, c, e) \in \mathcal{H}$ and define $(\tilde{\varphi}, \tilde{\mathfrak{m}}, \tilde{c}, \tilde{e}) := \Phi(\varphi, \mathfrak{m}, c, e)$. From the discussion in Section 3.1 and the definition of c_e , we obtain that $\tilde{\mathfrak{m}}$ is a p -hypermap. In order to see that $(\tilde{\varphi}, \tilde{\mathfrak{m}}, \tilde{c}, \tilde{e}) \in \mathcal{H}$, it thus remains to see that \tilde{e} is leaving \tilde{c}^+ in the orientation $\tilde{\varphi}$.

Similarly to Section 3.1, we denote the light side of the geodesic γ by $\ell = (\ell_1, \dots, \ell_k)$ and its dark side by $d = (d_1, \dots, d_k)$. Let also $c_{\tilde{e}}$ be the corner in $\tilde{\mathfrak{m}}$ that precedes \tilde{e} in the orientation $\tilde{\varphi}$. Everything follows from the following claim.

Claim. *In $\tilde{\mathfrak{m}}$, in the orientation $\tilde{\varphi}$, the lightest geodesic from $c_{\tilde{e}}$ to \tilde{c}^+ is the path $(\tilde{e}, d_{k-1} \times \ell_k, \dots, d_1 \times \ell_2)$.*

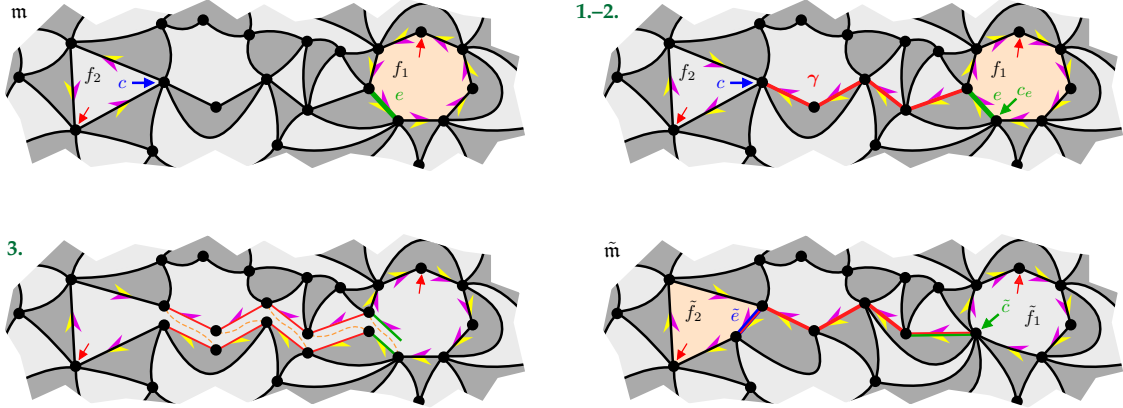


Figure 11: The involution $\Phi: \mathcal{H} \rightarrow \mathcal{H}$. Only the orientation around the faces of interest and along γ are depicted. **Top left.** The input. **Top right.** We changed the orientation and defined the sliding path γ . **Bottom left.** We slit along the path. The dashed lines indicate to sew back after sliding. **Bottom right.** The output.

Indeed, this implies that \tilde{e} is leaving \tilde{c}^+ in the orientation \tilde{o} , so that $(\tilde{o}, \tilde{m}, \tilde{c}, \tilde{e}) \in \mathcal{H}$. Furthermore, applying Φ to $(\tilde{o}, \tilde{m}, \tilde{c}, \tilde{e})$ means reverting \tilde{o} back to o and sliding back along the previous sliding path. More precisely, let us set $\tilde{\gamma} = (\tilde{e}, d_{k-1} \bowtie l_k, \dots, d_1 \bowtie l_2)$ and denote by $\tilde{\ell} = (\tilde{\ell}_1, \dots, \tilde{\ell}_k)$ and $\tilde{d} = (\tilde{d}_1, \dots, \tilde{d}_k)$ its light and dark sides. For $2 \leq i \leq k$, we thus have $\tilde{\ell}_{k+2-i} = l_i$ and $\tilde{d}_{k+1-i} = d_i$, so that we match back $\tilde{\ell}_{k+2-i} \bowtie \tilde{d}_{k+1-i} = l_i \bowtie d_i$. Furthermore, $\tilde{d}_k = d_1$ is unmatched, which amounts to matching it with the extra dangling edge l_1 that we suppressed. Finally, the distinguished elements correspond.

Let us then show the claim, using Figure 12 as support. We argue by contradiction and assume that the lightest geodesic $\tilde{\varphi}$ from \tilde{e} to \tilde{c}^+ differs from $\tilde{\gamma}$. Then $\tilde{\varphi}$ has to leave $\tilde{\gamma}$ at some point (either from its light side or from its dark side) and merge back at some other point (also either from its light side or from its dark side). This circumvention is either strictly shorter (a shortcut) or has the same length but comes before in the direction incoming edge, light face around the vertex where the splitting occurs. Tracing back this circumvention in m shows that this would have provided a better option than γ , creating either a shortcut or a better circumvention, thus contradicting the definition of γ . \square

Specialization. We now see how Φ specializes into a bijection interpreting (2). We let

$$\mathbf{a} = (a_1, \dots, a_r) \quad \text{and} \quad \tilde{\mathbf{a}} = (\tilde{a}_1, \dots, \tilde{a}_r) := (a_1 - 1, a_2 + 1, a_3, \dots, a_r)$$

be as in the statement of Proposition 1. Note that this means that p -hypermaps of type \mathbf{a} are either p -constellations or quasi- p -constellations whose **first** face is flawed. Similarly, p -hypermaps of type $\tilde{\mathbf{a}}$ are either p -constellations or quasi- p -constellations whose **second** face is flawed.

We fix an orientation o and define the following sets, whose cardinalities are respectively the left-hand side and the right-hand side of (2), by Proposition 3 (recall also the convention at the beginning of Section 2 for distinguishing corners).

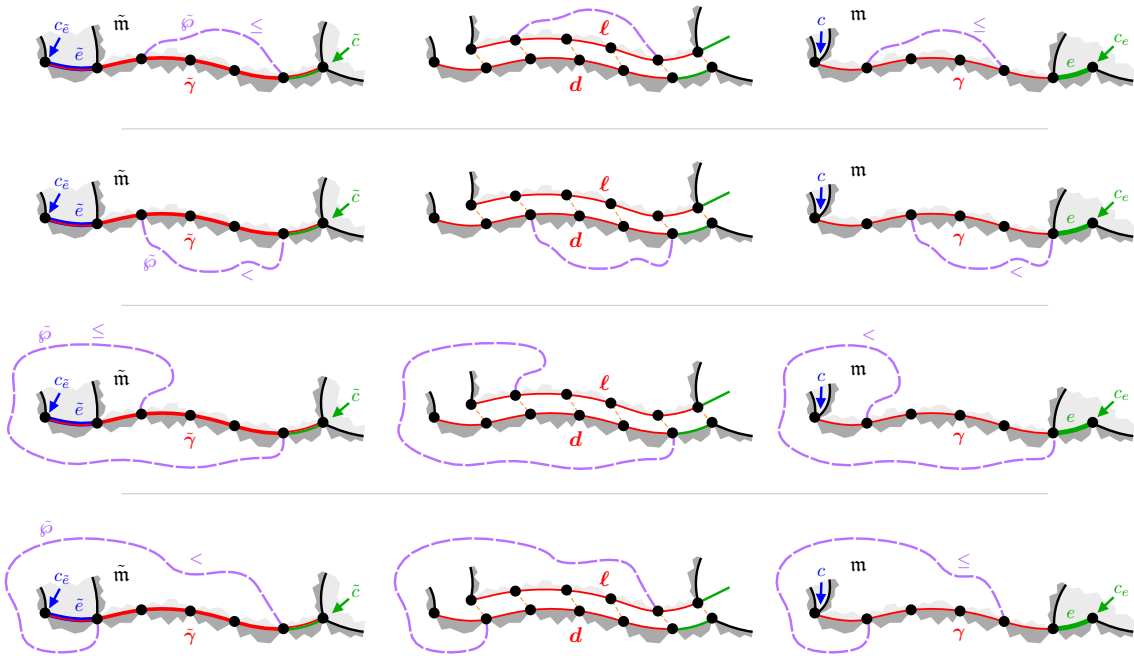
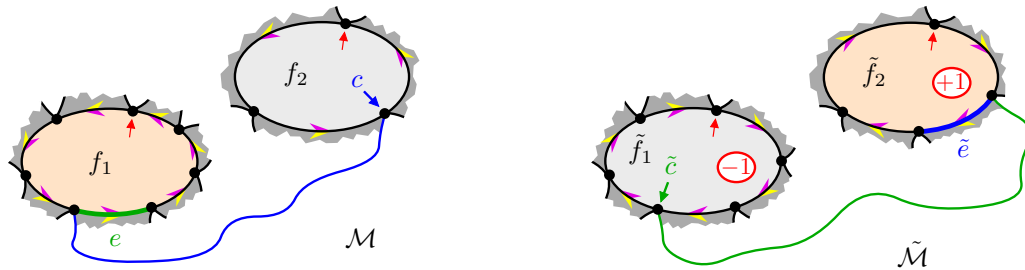


Figure 12: Proof of the fact that the sliding path γ in \mathfrak{m} used for applying Φ becomes the (reverse of the) sliding path in the image $\tilde{\mathfrak{m}}$ used for applying Φ a second time. **Left column.** If $\tilde{\varphi} \neq \tilde{\gamma}$, it means that $\tilde{\varphi}$ uses one of the purple dashed circumventions (each of the four lines corresponds to a possibility); if it starts on the light side, then it is shorter (symbol \leq); if it starts on the dark side, then it is strictly shorter (symbol $<$). **Middle and right column.** Tracking it back to the original map \mathfrak{m} yields a circumvention providing a better option than γ .

- ◇ We let \mathcal{M} be the set of p -hypermaps of type \mathfrak{a} carrying
 - one distinguished corner c in the **second** face,
 - one distinguished edge e incident to the **first** face and leaving c^+ , for the **orientation** \mathfrak{o} .

- ◇ We let $\tilde{\mathcal{M}}$ be the set of p -hypermaps of type $\tilde{\mathfrak{a}}$ carrying
 - one distinguished corner \tilde{c} in the **first** face,
 - one distinguished edge \tilde{e} incident to the **second** face and leaving \tilde{c}^+ , for the **reverse orientation** $\tilde{\mathfrak{o}}$.



Here, $p = 3$, we are in the case (ii) of Proposition 1, and $\mathfrak{o} = \text{light-left}$.

The pictograph above summarizes the definitions of \mathcal{M} and $\tilde{\mathcal{M}}$. The red arrowheads denote the marked corners of the depicted faces; the half-arrowheads along the edges symbolize the convened orientation. Finally, the red ± 1 on the right shows the increase or decrease of the degree of the face in $\tilde{\mathcal{M}}$ in comparison with the one of the corresponding face in \mathcal{M} . In order to avoid confusion, we denote the first and second faces of maps in \mathcal{M} by f_1 and f_2 as before, and use \tilde{f}_1 and \tilde{f}_2 instead, for maps in $\tilde{\mathcal{M}}$. The paths symbolize the fact that the edges are leaving the corners.

Remark 2. Note that the convention on the orientation of edges is not the same in the definitions of the sets \mathcal{M} and $\tilde{\mathcal{M}}$. This clearly bears no effects from an enumeration point of view but is of crucial importance for our bijections.

Corollary 5. *The mapping Φ specializes into a bijection from the set $\{(\mathfrak{o}, \mathfrak{m}, c, e) : (\mathfrak{m}, c, e) \in \mathcal{M}\}$ onto the set $\{(\tilde{\mathfrak{o}}, \tilde{\mathfrak{m}}, \tilde{c}, \tilde{e}) : (\tilde{\mathfrak{m}}, \tilde{c}, \tilde{e}) \in \tilde{\mathcal{M}}\}$, thus providing a bijection between \mathcal{M} and $\tilde{\mathcal{M}}$.*

3.3 Face of degree one

We proceed to the bijective interpretation for the identity (3) of Proposition 2, which works in a similar fashion as before.

Setting. Let $\mathbf{a} = (1, a_2, \dots, a_r)$ and $\tilde{\mathbf{a}} = (\tilde{a}_2, \dots, \tilde{a}_r) := (a_2 + 1, a_3, \dots, a_r)$ be tuples of positive integers, both with at most two coordinates not lying in $p\mathbb{N}$. In order not to be confused by the index shift in \tilde{a}_2 , we denote the faces of p -hypermaps of type $\tilde{\mathbf{a}}$ by $\tilde{f}_2, \dots, \tilde{f}_r$. In particular, p -hypermaps of type $\tilde{\mathbf{a}}$ are either p -constellations, or are quasi- p -constellations whose face \tilde{f}_2 (the one with degree \tilde{a}_2) is flawed. We fix an orientation \mathfrak{o} and define the following sets, whose cardinalities are the sides of (3), again by Proposition 3 for the right-hand side.

- ◇ We let \mathcal{N} be the set of p -hypermaps of type \mathbf{a} carrying
 - one distinguished corner c in the face f_2 .
- ◇ We let $\tilde{\mathcal{N}}$ be the set of p -hypermaps of type $\tilde{\mathbf{a}}$ carrying
 - one distinguished vertex \tilde{v} ,
 - one distinguished edge \tilde{e} incident to \tilde{f}_2 and leaving \tilde{v} for the orientation $\tilde{\mathfrak{o}}$.



We put f_1 on the pictograph since we think of it as the “missing” distinguished element for \mathcal{N} . Note that we do not need to specify an orientation for maps in \mathcal{N} ; we will however use the orientation \mathfrak{o} for these maps in due time. The bijections between \mathcal{N} and $\tilde{\mathcal{N}}$ can be thought of as degenerate versions of the one of the previous section. Here, we do not have an involution; we need to describe both mappings. We break them down into similar steps as above. See Figure 13.

Suppressing a face. We consider $(m, c) \in \mathcal{N}$.

1. Reorientation

From this point on, we use the reverse orientation \tilde{o} .

2. Sliding path

We forget the mark of the unique corner of f_1 and consider the lightest geodesic γ from this corner to c^+ .

3. Slitting, sliding, sewing

We denote by d_0 the unique edge incident to f_1 . We slit, slide, sew along γ from the unique corner of f_1 to c as described in Section 3.1, while furthermore matching the unmatched light side of the first edge with d_0 , that is, $\ell_1 \bowtie d_0$ in the notation of Section 3.1.

4. Output

We set $\Psi_-(m, c) := (\tilde{m}, \tilde{v}, \tilde{e})$, where \tilde{m} is the resulting map, \tilde{e} is the edge corresponding to the unmatched dark side of the final edge of γ , and \tilde{v} is the origin of γ , that is, $\tilde{v} := (\ell_1 \bowtie d_0)^-$.

Note that the face f_1 has been suppressed in the process.

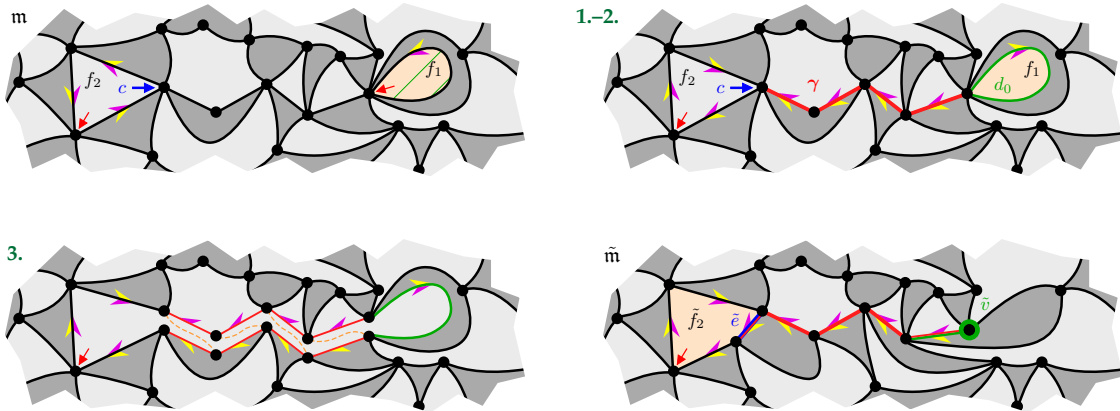


Figure 13: The bijection in the case of a degree 1-face, from \mathcal{N} to $\tilde{\mathcal{N}}$.

Adding a face. We consider $(\tilde{m}, \tilde{v}, \tilde{e}) \in \tilde{\mathcal{N}}$.

1. Reorientation

We reverse the orientation, which means that we use the original orientation o .

2. Sliding path

We consider the corner $c_{\tilde{e}}$ preceding \tilde{e} and the lightest geodesic $\tilde{\gamma}$ from $c_{\tilde{e}}$ to \tilde{v} .

3. Slitting, sliding, sewing

We slit \tilde{m} along $\tilde{\gamma}$, entering from $c_{\tilde{e}}$ and stopping at \tilde{v} , **without disconnecting the map at \tilde{v} .**

Denoting as before by $\tilde{\ell} = (\tilde{\ell}_1, \dots, \tilde{\ell}_k)$ and $\tilde{\mathbf{d}} = (\tilde{d}_1, \dots, \tilde{d}_k)$ the light and dark sides of $\tilde{\gamma}$, this means that $\tilde{\ell}_k^+ = \tilde{d}_k^+$.

We then sew back ℓ onto \mathbf{d} as before, by matching ℓ_{i+1} with d_i , for every $1 \leq i \leq k-1$. In this matching, $\tilde{\ell}_k^+$ is matched with \tilde{d}_k^- , so that \tilde{d}_k creates a loop, whose unmatched side encloses an extra face, which we denote by f_1 and mark at its unique corner.

4. Output

As usual, we replace the unmatched $\tilde{\ell}_1$ with a corner, which we denote by c . We let \mathfrak{m} be the resulting map, and set $\Psi_+(\tilde{\mathfrak{m}}, \tilde{v}, \tilde{e}) := (\mathfrak{m}, c)$.

Theorem 6. *The mappings $\Psi_-: \mathcal{N} \rightarrow \tilde{\mathcal{N}}$ and $\Psi_+: \tilde{\mathcal{N}} \rightarrow \mathcal{N}$ are well defined and inverse bijections.*

Proof. The proof closely follows that of Theorem 4 and is left to the reader. \square

4 Uniform sampling

The bijections presented above can be used to sample a uniform p -constellation or a quasi- p -constellation of a given type $\mathbf{a} = (a_1, \dots, a_r)$ as follows. Without loss of generality, let us assume that p divides every a_i for $1 \leq i \leq r-2$. In order to address this question, we first denote by $n := D(\mathbf{a})$ the number of dark faces and define the finite sequence of tuples:

$$(np), \quad (1, np-1), \dots, (a_1, np-a_1), \quad (a_1, 1, np-a_1-1), \dots, (a_1, a_2, np-a_1-a_2), \\ (a_1, a_2, 1, np-a_1-a_2-1), \dots, (a_1, a_2, \dots, a_r)$$

which we denote by $\mathbf{a}^0, \dots, \mathbf{a}^{np-a_r} = \mathbf{a}$ for simplicity. The important point is that each tuple of the sequence has at most two coordinates not lying in $p\mathbb{N}$ and, for each $1 \leq j \leq np-a_r$, the tuple \mathbf{a}^j differs from \mathbf{a}^{j-1} by -1 in one coordinate and

- (a) either $+1$ in another coordinate,
- (b) or by a new coordinate equal to 1 ,

while all unchanged coordinates being multiples of p . In particular, up to an obvious coordinate permutation, this change in coordinates is covered by Propositions 1.(i) or 2. We chose this particular sequence but, in fact, any other sequence with the above properties also works. Note also that the number of dark faces corresponding to any of the tuples in the sequence is equal to n .

We sample a sequence of p -hypermaps $\mathfrak{m}^0, \dots, \mathfrak{m}^{np-a_r}$ such that, for every $1 \leq j \leq np-a_r$, the p -hypermap \mathfrak{m}^j is uniformly distributed among p -hypermaps of type \mathbf{a}^j as follows.

- ◇ First of all, using the simple bijection of Figure 3, sampling a uniform p -constellation \mathfrak{m}^0 of type $\mathbf{a}^0 = (np)$ amounts to sampling a uniform p -ary tree with n nodes, which can be done for instance by sampling an appropriate Łukasiewicz path.
- ◇ Assuming \mathfrak{m}^{j-1} has been sampled, we sample \mathfrak{m}^j as follows.
 - In \mathfrak{m}^{j-1} , we distinguish either a corner or a vertex:
 - (a) if a coordinate of \mathbf{a}^{j-1} is increased by $+1$, we distinguish a corner chosen uniformly at random in the corresponding face of \mathfrak{m}^{j-1} ;

- (b) otherwise, we distinguish a vertex chosen uniformly at random in m^{j-1} .
- Then, we distinguish an edge chosen uniformly at random among those incident to the face corresponding to the decreasing coordinate of a^{j-1} that leave the distinguished corner or vertex.
- We apply the mapping of Section 3.2 in case (a) or that of Section 3.3 in case (b) in order to sample a p -hypermap m^j of type a^j from m^{j-1} , the distinguished corner or vertex, and the distinguished edge.
- In the resulting p -hypermap m^j , we forget the distinguished elements coming from the bijection.

As the number of ways to select the desired distinguished elements in m^{j-1} only depends on the type a^{j-1} , distinguishing elements does not bias the uniform probability. And since bijections preserve the uniform distribution, the fact that m^{j-1} is uniformly distributed among the p -hypermaps of type a^{j-1} implies that m^j is uniformly distributed among the p -hypermaps of type a^j . By induction, we obtain that, for every $1 \leq j \leq np - a_r$, the p -hypermap m^j is indeed uniformly distributed among p -hypermaps of type a^j . In particular, m^{np-a_r} is a uniform p -constellation or a quasi- p -constellation of type a , as desired.

Remark 3. Seeing the last face of each p -hypermap in the sequence as an *external face*, we obtain a sequence of “growing” uniform p -hypermaps where the faces are added one by one, each face being grown by adding one degree at each step. In other words, we start from a uniform tree made of n dark degree p -faces. Then we grow a first face of degree 1, then 2, and so on until it has degree a_1 , then we grow a second face of degree 1, then 2, and so on until it has degree a_2 , and we keep growing faces until we reach the desired type a .

References

- [Bet14] Jérémie Bettinelli. Increasing forests and quadrangulations via a bijective approach. *J. Combin. Theory Ser. A*, 122(0):107–125, 2014.
- [Bet19] Jérémie Bettinelli. Slit-slide-sew bijections for bipartite and quasibipartite plane maps. In *31th International Conference on Formal Power Series and Algebraic Combinatorics (FPSAC 2019)*, volume 82B of *Sém. Lothar. Combin.*, pages Art. 82, 1–12. SLC, 2019.
- [Bet20] Jérémie Bettinelli. Slit-slide-sew bijections for bipartite and quasibipartite plane maps. *Electron. J. Combin.*, 27(3):Paper 3.4, 23, 2020.
- [BMS00] Mireille Bousquet-Mélou and Gilles Schaeffer. Enumeration of planar constellations. *Adv. in Appl. Math.*, 24(4):337–368, 2000.
- [CF14] Gwendal Collet and Éric Fusy. A simple formula for the series of constellations and quasi-constellations with boundaries. *Electron. J. Combin.*, 21(2):Paper 2.9, 27, 2014.
- [GJ83] Ian P. Goulden and David M. Jackson. *Combinatorial enumeration*. Wiley-Interscience Series in Discrete Mathematics. John Wiley & Sons, Inc., New York, 1983. With a foreword by Gian-Carlo Rota, A Wiley-Interscience Publication.
- [Tut62] William T. Tutte. A census of slicings. *Canad. J. Math.*, 14:708–722, 1962.