Journal of Knot Theory and Its Ramifications
Vol. 15, No. 8 (2006) 1001–1035

BLOBS AND FLIPS ON GEMS

SÓSTENES LINS
Depto. de Matematica, Univ. Federal de Pernambuco,
Recife, Brazil
sostenes.lins@gmail.com

MICHELE MULAZZANI
Dip. di Matematica, Univ. di Bologna, Bologna, Italy
mulazzani@dm.unibo.it

Accepted 24 June 2006

ABSTRACT

In this paper we prove that two n-gems induce the same manifold if and only if they are linked by a finite sequence of *gem moves*. A gem move is either a *blob move*, consisting in the creation or cancellation of an n-dipole, or a *clean flip*, which is a switch of a pair of edges of the same color that thickens an h-dipole, $1 \leq h \leq n-1$, or the inverse operation, which slims an h-dipole, $2 \leq h \leq n$. Moreover we prove that we can reorder the gem moves, so that all the blob creations precede all clean flips which then precede all the blob cancellations. This reordering is of interest because it is an easy matter to decide whether two gems are linked by a finite sequence of clean flips. As a consequence, if a bound for the number of blob creations is established, then there exists a deterministic finite algorithm to decide whether two gems induce the same manifold or not.

Keywords: Colored pseudo-triangulation; gem; dipole move; barycentric subdivision and thickening; Pachner move; Gagliardi bisection; blob; flip; clean flip; gem move; permutohedron.

Mathematical Subject Classification 2000: Primary 57Q05, 57M15; Secondary 57M12, 57M25, 05C10

1. Introduction

Gem theory, via its dual, the theory of colored pseudo-triangulations of PL-manifolds, can be viewed as an intermediate theory between that of simplicial triangulations, where the number of simplices is large and the pasting description is trivial, and that of CW-complexes, where we have few simplices, but the pasting is complex, making it unsuitable for a combinatorial description. In dimension three a point in which gem theory differs from other descriptions of 3-manifolds is that the set of minimal forms act as an attractor under a rich combinatorial simplification theory. This makes 3-gems an adequate theory for the recognition and the computational classification of "small" 3-manifolds (see [10]). A very attractive aspect of

3-dimensional gem theory is its direct connection with the theory of framed links. In fact, it is possible to obtain a gem from a *blink* (blackboard framed link) by replacing each crossing of the blink by a subgem with 12 vertices (see [7, Sec. 13.5]). This result can be further improved by decreasing the number of vertices which replaces each crossing from 12 to 8. Another basic result is a direct construction, in terms of 3-gems, of the whole class of closed 3-manifolds in terms of "twistors". A *twistor* in a gem is a special kind of embedded solid torus induced by a special pair of vertices [11].

Beyond $n = 3$, it has been an open problem to produce the barycentric subdivision of a colored pseudo-triangulation of an n-manifold (dual of an n-gem) by displaying a finite sequence of dipole moves. In this paper we exhibit such a sequence for all cases $n \geq 1$. As a consequence, we prove the enhanced theorem stated below and proved as Proposition 5. A *blob* in an n-gem is a K-dipole with $|K| = n$. For a non-negative α, let G^α be the gem obtained from gem G by creating α blobs at arbitrary edges of G. A *flip* at a pair of equally colored edges in a gem is the interchange of the two edges by a new pair, having the same ends and color. We distinguish a class of flips, named *clean flips*, which maintain the induced PL n-manifold. A blob *creation* or *cancellation* is the creation/cancellation of the K-dipole and, therefore, maintains the induced PL-manifold. Thus, if we want to test whether two n-gems induce the same manifold, by creating blobs in the gem with less vertices, there is no loss of generality in assuming that they have the same number of vertices. We prove the following enhanced equivalence theorem: if G and H are gems with the same number of vertices inducing the same n-manifold, then there is an integer $\alpha(G, H)$ such that G^α and H^α are linked by a finite number of clean flips, where G^α and H^α are respectively G and H with $\alpha = \alpha(G, H)$ blobs put over arbitrary edges. Consequently, it is enough to establish a bound for $\alpha(G, H)$ in order to produce an algorithm that decides whether G and H induce or not the same n-manifold. Our result has, as a corollary (providing an independent proof), a basic consequence of the Ferri–Gagliardi theorem on crystallizations [4], stating that two n-gems induce the same n-manifold if and only if they are linked by dipole moves.

2. Formal Definitions and the Weak Equivalence Theorem

For the basic notions on PL topology, we refer to [15]. An *n-pseudogem* G is a finite $(n+1)$-regular edge colored graph, with color set $N = \{0, 1, \ldots, n\}$, such that at each vertex the incident edges have distinct colors. The vertex set of G is denoted by $V(G)$. For $\emptyset \subset K \subset N$, the *$K$-graph* of G is the subgraph of G induced by the k-colored edges (or *k-edges*), with $k \in K$. A *K-residue of G* is any connected component of the K-graph of G. If $v \in V(G)$, the K-residue of G containing v is denoted by v_K^G, and v is said to *represent* the residue. A *K-subgraph of G* is the union of a non-null set of K-residues of G. A K-residue of G and the K-graph of G are extremal examples of K-subgraphs of G. Note that a K-subgraph of G is itself a $(|K| - 1)$-gem whose set of colors is $K \subseteq N$.

Each n-pseudogem G induces a topological space $|G|$, which is an n-dimensional pseudo-manifold (see [5]). An *n-gem* is an n-pseudogem inducing an n-manifold. This happens if and only if each J-residue induces a $(|J|-1)$-sphere. All PL n-manifolds admit gem descriptions. Henceforth we restrict the theory to gems, even though some aspects generalize to pseudogems. A basic property of gems is that an n PL-manifold is orientable if and only if any gem inducing it is a bipartite graph [9].

A gem G is the 1-skeleton of the dual ${}^dT$ of a colored pseudo-triangulation T of the manifold $|G|$, which is endowed with a coloring on its 0-simplices so that the vertices of each simplex are differently colored. This coloring is induced by the $(n+1)$-coloring of the edges of G: a j-edge $(i = 0, 1, \ldots, n)$ of G intersects an $(n-1)$-simplex S of T in a single point. There are exactly two n-simplices, say S' and S'', containing S. Color the 0-simplex of S' (respectively S'') that is not in S with the color j. On the other hand, the cellular complex ${}^dT$ can be obtained from G in the following way: attach a disk to each K-residue, with $|K| = 2$; then attach a 3-ball to each K-residue, with $|K| = 3$; and so on, until the attaching of an n-ball to each K-residue, with $|K| = n$.

For $\emptyset \subset K \subset N$, a *$K$-dipole* in an n-gem G is a K-residue containing only two vertices, say x and y, such that $x^G_{N\backslash K} \neq y^G_{N\backslash K}$. By duality, we can also talk about dipoles in colored pseudo-triangulations. The *cancellation* of a K-dipole with vertices x, y in a gem G is the following operation. Remove the edges and vertices of the dipole and weld the pendant edges with the same color (there is a pair of these for each color in $N\backslash K$). The *creation* of a K-dipole is the inverse operation. A *dipole move* is either a dipole creation or a dipole cancellation. By duality, dipole moves also apply to colored pseudo-triangulations. Dipole moves on gems (as well as in their duals) do not change the homeomorphism type of the induced manifolds.

In this paper we constructively obtain a sequence of dipole moves linking a colored pseudo-triangulation of an n-manifold to its barycentric subdivision. This result was previously unknown for dimensions $n > 3$. Together with Casali's work [1], it gives a simple and clearer proof of the sufficiency of the dipole moves for the homeomorphism problem of PL n-manifolds. In fact, we define and work with *gem moves* (shortly to be defined) which, in our context, are better to work with than dipole moves. In doing so, we obtain our main result (Proposition 5), which is an enhanced form of the Ferri–Gagliardi equivalence theorem [4].

The *barycentric thickening of an n-gem* G is the dual of the barycentric subdivision of the colored pseudo-triangulation dual of G. Consider the following streamlined proof of the weak equivalence theorem for gems. The adjective weak refers to the use of the barycentric thickenings and their inverses. If two gems G and H are linked by a finite sequence of dipole moves, then we write $G \overset{*}{\longleftrightarrow} H$.

Proposition 1 (Weak equivalence theorem). *Two gems G_1 and G_2 induce the same n-manifold if and only if they are linked by barycentric thickenings, dipole moves and inverse barycentric thickenings.*

Proof. In the diagram below, $G_i^\star, G_i^{\star\star}$ $(i = 1, 2)$ are the first and the second barycentric thickenings of G_i. The meaning of $H_i \rightleftharpoons H_{i+1}$ is that ${}^dH_{i+1}$, the dual of H_{i+1}, is obtainable from ${}^dH_i$, the dual of H_i, by a single Pachner move. Even though H_i might not be a gem, except for $i = 0$ and $i = p$, $H_i^\star$ is a gem for $i = 0, 1, \ldots, p$.

$$\begin{array}{ccccccccc} G_1 \to & G_1^\star = H_0 & \rightleftharpoons & H_1 & \rightleftharpoons & \cdots & \rightleftharpoons & H_p = G_2^\star & \leftarrow G_2 \\ & \downarrow & & & & & & \downarrow & \\ & G_1^{\star\star} = H_0^\star & \xleftrightarrow{*} & H_1^\star & \xleftrightarrow{*} & \cdots & \xleftrightarrow{*} & H_p^\star = G_2^{\star\star} & \end{array}$$

Note that ${}^dH_0$ and ${}^dH_p$ are simplicial complexes. So they induce the same n-manifold if and only if they are linked by a finite sequence of Pachner moves [16, 8]. Moreover, we claim that ${}^dH_i$ and ${}^dH_{i+1}$, are linked by a Pachner move if and only if $H_i^\star$ and $H_{i+1}^\star$, the duals of the barycentric subdivisions of ${}^dH_i$ and ${}^dH_{i+1}$ are linked by a finite sequence of dipole moves. This result first appears in [1, Lemma 5] but, for completeness, we provide a simplified proof of this fact. Let ${}^dH_{i+1}$ be obtained from ${}^dH_i$ by a single k-bistellar move $\chi(A, B)$, with $0 \leq k \leq n$, being $k = \dim A$ and $n - k = \dim B$. This implies, according to Pachner's theory, that they differ only in a subcomplex C, which is triangulated as the join complex $\partial A * B$ in ${}^dH_i$ and as the join complex $A * \partial B$ in ${}^dH_{i+1}$. Note that the boundary of C is $\partial A * \partial B$ and remains invariant under this replacement. In order to prove that $H_{i+1}^\star$ can be obtained from $H_i^\star$ by a finite number of dipole moves, it suffices to prove that $H_i^\star$ is dipole equivalent to H, where ${}^dH$ is obtained from ${}^dH_i^\star$ by replacing the barycentric subdivision of $C = A * \partial B$ with the cone over the barycentric subdivision of ∂C. By interchanging A and B, as well as i and $i+1$, the same proof applies and sufficiency is established.

Observe that $C = A * \partial B$ consists of $n - k + 1$ n-simplices $\sigma_1, \ldots, \sigma_{n-k+1}$. Let s_{ij} be the common $(n-1)$-face of σ_i and σ_j, for $1 \leq i < j \leq n - k + 1$. Let S_1 be the $(n-1)$-skeleton of C and let $S' = \overline{S_1 \backslash \partial C}$. Obviously, S' is exactly the union of the s_{ij}'s previously defined. Note that the triangulation T_1 of $C^\star$ consists of the union of the cones over the barycentric subdivision of the boundary of any n-simplex of C (i.e. the components of $C \backslash S_1$). Now, let us consider the n-simplices σ_1 and σ_2. In $C^\star$, the simplex s_{12} is subdivided into $n!$ $(n-1)$-simplices $\sigma_1', \ldots, \sigma_{n!}'$, each being the face of one n-simplex in the barycentric subdivision of σ_1, and the face of one n-simplex in the barycentric subdivision of σ_2. We can suppose, up to reordering, that σ_h' has at least an $(n-2)$-face in common with $\sigma_1' \cup \cdots \cup \sigma_{h-1}'$, for any $h = 2, \ldots, n!$. By performing the sequence of $n!$ dipole cancellations (all involving color n), corresponding to the pairs of n-simplices with common face $\sigma_1', \ldots, \sigma_{n!}'$ respectively, we obtain a new triangulation T_2 of $C^\star$, which consists of the union of the cones over the barycentric subdivision of the boundary of any component of $C \backslash S_2$, where $S_2 = \overline{S_1 \backslash s_{12}}$. If $k = n - 1$, we have only one component, since in this case $S' = s_{12}$, and the statement is obtained. If $k < n - 1$, then consider the n-simplex σ_3. The $(n-1)$-simplex s_{13} (respectively, s_{23}) is subdivided in $n!$ $(n-1)$-simplices $\sigma_1'', \ldots, \sigma_{n!}''$ (respectively, $\sigma_1''', \ldots, \sigma_{n!}'''$), each being the face of one n-simplex in the barycentric subdivision of σ_3 and the face of one n-simplex in the subdivision of $\sigma_1 \cup \sigma_2$ in the triangulation T_2 of $C^\star$. Again we can suppose, up to reordering, that σ''_h (respectively, σ'''_h) has at least one $(n-2)$-face in common with $\sigma''_1 \cup \cdots \cup \sigma''_{h-1}$ (respectively, $\sigma'''_1 \cup \cdots \cup \sigma'''_{h-1}$), for any $h = 2, \ldots, n!$. By performing the sequence of $2n!$ dipole cancellations corresponding to these pairs of n-simplices with common face $\sigma''_1, \ldots, \sigma''_{n!}, \sigma'''_1, \ldots, \sigma'''_{n!}$ respectively, we obtain a new triangulation T_3 of $C^\star$, which consists of the union of the cones over the barycentric subdivision of the boundary of any component of $C \backslash S_3$, where $S_3 = \overline{S_2 \backslash s_{13} \cup s_{23}}$. If $k = n-2$, we have only one component, since in this case $S' = s_{12} \cup s_{13} \cup s_{23}$, and the statement is obtained. If $k < n-2$, we proceed by induction. At the end of the process, consisting of $n-k$ steps, we get the final triangulation of $C^\star$, which is the cone over the barycentric subdivision of the boundary of C. This completes the proof of the claim. Thus, G_1 and G_2 induce the same manifold if and only if G_1 is linked to G_2 by two barycentric thickenings, a finite sequence of dipole moves and two inverse barycentric thickenings. □

A central objective of this paper is to show that the above theorem remains true for all n if we leave out the barycentric thickenings and their inverses. This is known for $n \leq 3$, but the case $n \geq 4$ is first proved here.

3. Blobs, Flips, Gem Moves and the Main Result

A *z-blob* is an $(N\backslash\{z\})$-dipole. The cancellation of a z-blob with vertices u', w' produces a z-edge $e \equiv (u, w)$. We say that the inverse operation is the *creation of a blob over the z-edge* (u, w). See Fig. 1. After such a creation there appears a pair of new z-edges (u, u') and (v', v). A *blob* is a z-blob for some $z \in N$. A *blob move* is either the creation or the cancellation of a blob. In the pictures throughout the article, a thick edge with a label J incident to a vertex v (represented by a small disk) means a set of $|J|$ edges, each one incident to v, colored by the elements of J. Thus, a thick edge labelled J incident to a pair of vertices v and w means a set of $|J|$ *parallel edges*, i.e. all with the same ends v and w. Since the a blob at a z-edge has its color set $N\backslash\{z\}$, the blob is completely specified by a circle over the edge. For this reason, we say that the blob is *over* the edge. Throughout the pictures we show two kinds of vertices, represented by black and white discs. Since we only work with gems, each n-residue induces an $(n-1)$-sphere and so is bipartite. The black/white partition is the manifestation of the bipartiteness of the n-residue, which, in the required cases, contains the partial gem being depicted. Thus we do not effect a flip like the one shown in the third subgraph of Fig. 2, which destroys the bipartiteness of an n-residue.

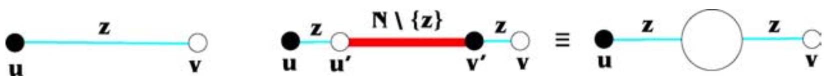

Fig. 1. Blob moves: creation or cancellation of a blob.

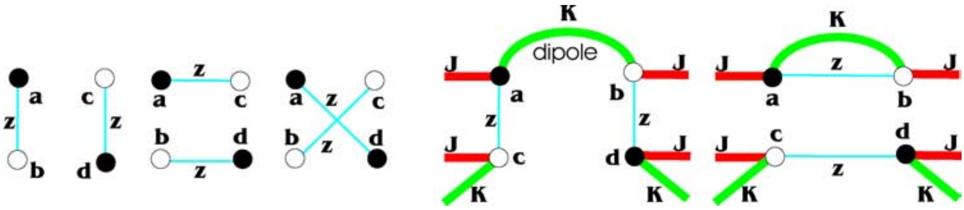

Fig. 2. Flips and a clean flip.

Let e and f be z-edges of an n-gem with ends (a, b) and (c, d) respectively. A *z-flip* is the operation which replaces e and f, either by z-edges e' with ends (a, c) and f' with ends (b, d), or else by z-edges e'' with ends (a, d) and f'' with ends (b, c). So, given two z-edges e and f, there are two possibilities of z-flips. A flip is a z-flip for a certain color z. There is a particular case of flip which is important for us. Assume that e and f are z-edges which are incident respectively to the vertices a and b of a K-dipole D, with $z \notin K$. Consider the flip replacing e and f by e' and f' respectively, so that the ends of e' are a and b. Such a flip, which *thickens* the dipole D to $D \cup \{e'\}$ is called a *t-flip*. The inverse operation, which *slims* the dipole $D \cup \{e'\}$ to D, is called an *s-flip*. A *clean* flip is either a t-flip or an s-flip. We shall see that, when applied to gems, a clean flip does not introduce singularities in the associated n-manifolds, as is the case of an arbitrary flip. More strongly, by Proposition 3, a clean flip leaves invariant the induced manifold.

Proposition 2 (Dipole thickening flips: t-flips). *Let a, b, e, f, e', f', D, K as in the definition of a t-flip in a gem G, and let H be the gem after the t-flip. Then $D \cup \{e'\}$ is a $(K \cup \{z\})$-dipole in H.*

Proof. Just note that $a^H_{N\backslash(\{z\}\cup K)} \neq b^H_{N\backslash(\{z\}\cup K)}$, since $a^G_{N\backslash K} \neq b^G_{N\backslash K}$. $\square$

Proposition 3 (Dipole factorization of a clean flip). *Let H be obtained from an n-gem F by an s-flip. Then there exists an n-gem G, obtained from F by a single dipole creation and producing H by a single dipole cancellation.*

Proof. Let $N = \{z\} \cup J \cup K$. We refer to Fig. 3, where $a^F_{\{z\}\cup K}$ and a^H_K are dipoles. Since a^H_K is a dipole, so is a^G_K. Indeed, $a^H_{\{z\}\cup J}$ and $a^G_{\{z\}\cup J}$ are isomorphic, so $a^G_{\{z\}\cup J}$ and $b^G_{\{z\}\cup J}$ are distinct. From F to G we have created the dipole b^G_K and from G to H we have cancelled the dipole b^G_J (or the dipole $c^G_{K\cup\{z\}}$). $\square$

Note that a flip maintains the set of vertices of an n-gem. Since dipole moves in n-gems maintain the induced n-manifold, according to Proposition 3, two n-gems linked by a finite number of clean flips have precisely the same set of vertices and induce the same n-manifold. If two n-gems G_1 and G_2 are linked by a single clean flip we write $G_1 \square G_2$. The notation $G_1 \square^r G_2$ means that we can go from G_1 to G_2

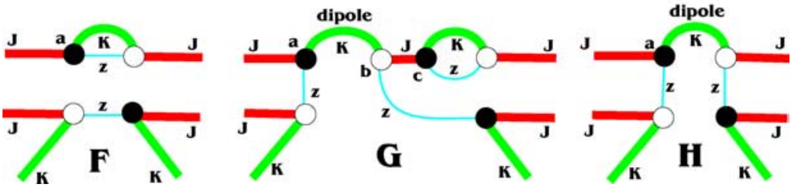

Fig. 3. Dipole factorization of an s-flip.

by effecting exactly r clean flips, and $G_1 \Box^* G_2$ means that the passage is effected by a finite (but not specified) number of clean flips.

Proposition 4 (Blob rotation). *Let the n-gems G_h and G_z differ only by the positioning of blobs B_h and B_z. Assume that B_h is over an h-edge incident to a vertex a, and that B_z is over the z-edge incident to a. Then $G_1 \Box^2 G_2$.*

Proof. Let b be the neighbor of a by color h and $K = N \backslash \{h, z\}$. We refer to Fig. 4. The s-flip which slims the dipole $b_{K \cup \{z\}}$ by a z-edge, followed by the t-flip which thickens the dipole b_K by an h-edge, perform the rotation of the blob. Note that the rotation maintains the set of vertices, namely $\{b, c\}$, of the blob. $\Box$

Given a connected n-gem G and a non-negative integer α, let G^α be an n-gem obtained from G by creating α blobs at arbitrary positions. An iterated application of the previous result shows that, modulo clean flips, the specific positions of these blobs are irrelevant.

Proposition 5 (Main result: enhanced equivalence theorem). *If G and H are n-gems inducing the same n-manifold, $|V(G)| \leq |V(H)|$, then there is an integer $\alpha = \alpha(G, H)$ such that $G^{\alpha'} \ \Box^* \ H^\alpha$, where $\alpha' = \alpha + (|V(H)| - |V(G)|)/2$.*

In order to prove the result, we state in Proposition 9 another result which implies it, in terms of gem moves. A *gem move* in a gem is either the creation/cancellation of a blob or a clean flip. If H is obtained from G by a single blob creation (cancellation) we write $G \uparrow H$ (respectively, $G \downarrow H$). If H is obtained

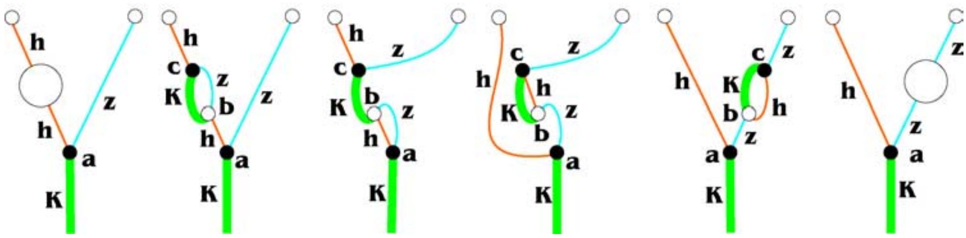

Fig. 4. Blob rotation.

from G by a finite number of blob creations (cancellations) we write $G \uparrow^* H$ (respectively, $G \downarrow^* H$). The notation $G(\uparrow \square \downarrow)H$ means that H is obtained from G by a single gem move. $G(\uparrow \square \downarrow)^* H$ means that H is obtained from G by a finite number of gem moves. In view of Proposition 6, we prove our main results using gem moves, which are easier to work with than dipole moves.

Proposition 6 (Gem moves $\times$ dipole moves). *Let G and H be n-gems. Then*

$$(G \overset{*}{\longleftrightarrow} H) \Leftrightarrow G(\uparrow \square \downarrow)^* \ H.$$

Proof. The creation and cancellation of a blob are dipole moves. A clean flip is factorable as a pair of dipole moves, as shown in Proposition 3. Reciprocally, by Proposition 2, the cancellation of a dipole can be accomplished by iterated t-flips at the dipole, until a blob is created. This blob is then cancelled. Thus, a dipole cancellation is factored as a finite sequence of gem-moves. The inverse sequence creates the dipole. $\square$

Proposition 7 (Commuting blob moves and clean flips). *Let G_1, G_2, and G_3 be n-gems. The following implications hold*:

$$(G_1 \downarrow G_2 \square G_3) \Rightarrow \exists\ G_2' \mid (G_1 \square^* G_2' \downarrow G_3),$$
$$(G_1 \square\ G_2 \uparrow G_3) \Rightarrow \exists\ G_2' \mid (G_1 \ \uparrow G_2' \square^* G_3).$$

Proof. In the first implication we can put (if necessary by clean flips) the blob which needs to be cancelled over another edge, so as not to interfere with the clean flip to be performed. Then we can perform the clean flip, obtaining G_2', and cancel the blob in its new location. This establishes the first implication. To prove the second, we start by creating a blob over an edge which does not interfere with the clean flip to be performed, thus defining G_2'. By clean flips, bring the created blob back to its appropriate location. This establishes the second implication. $\square$

Proposition 8 (Blob conjugation). $[G(\uparrow \square \downarrow)^* H] \Leftrightarrow \exists\ G', H' \mid [G \uparrow^* G' \square^* H' \downarrow^* H]$.

Proof. The implication $[G(\uparrow \square \downarrow)^* H] \Leftarrow \exists\ G', H' \mid [G \uparrow^* G' \square^* H' \downarrow^* H]$ is obvious. To prove the reverse implication, apply the second implication of Proposition 7 as many times as necessary, so as to get all the creations of blobs before any clean flip. This defines G'. Next apply the first implication of Proposition 7 as many times as necessary, so as to have all the blob cancellations after any clean flip. This defines H'. Cancellations of blobs yields H. $\square$

4. Equivalence Theorem, Bisections and Trisections

We are now in a position to state a result (the equivalence theorem in terms of gem moves) which implies our main result (Proposition 5).

Proposition 9 (Equivalence theorem). *Two gems G and H induce the same n-manifold if and only if they are linked by a finite number of gem moves, that is,* $G\ (\uparrow \Box \downarrow)^*\ H$.

Equivalence theorem $\Rightarrow$ main result. Assume that the equivalence theorem holds, and that G and H are n-gems under the hypothesis of the main result, Proposition 5. It follows that $G(\uparrow \Box \downarrow)^* H$. By Proposition 8, there exists two gems G' and H' such that $G \uparrow^* G' \Box^* G'' \downarrow^* H$. This means that there exists an $\alpha(G, H)$ as stated in the main result, thereby proving the implication. $\Box$

Proposition 10 (Lemma on barycentric thickening). *Any n-gem G and its barycentric thickening $G^\star$ are linked by a finite number of gem moves, that is,* $G\ (\uparrow \Box \downarrow)^*\ G^\star$.

Lemma on barycentric thickening $\Rightarrow$ equivalence theorem. Assume that the lemma on barycentric thickening holds. Then, by Proposition 1, we have $G \overset{*}{\longleftrightarrow} G^*$, which, by Proposition 6, is equivalent to $G\ (\uparrow \Box \downarrow)^*\ G^\star$. $\Box$

Proof of the lemma on barycentric thickening. We show below, in the remaining part of this section, that the passage $G \longrightarrow G^\star$ can be factored as

$$G = G_0 \ \rightarrow_t \ G_1 \ \rightarrow_t \ G_2 \ \rightarrow_t \ \cdots \ \rightarrow_t \ G_{q-1} \ \rightarrow_t \ G_q = G^\star,$$

where each step $G_{\ell-1} \rightarrow_t G_\ell$, $0 \leq \ell \leq q-1$, is effected by an operation called *trisection*, defined below. Along Secs. 6–8 we prove that $G_{\ell-1} \rightarrow_t G_\ell$ can be accomplished by gem moves, namely, $G_{\ell-1}(\uparrow \Box \downarrow)^*\ G_\ell$. Thus, up to the definition of trisection, the proof that a barycentric thickening can be factored by trisections and the proof that a trisection is accomplished by gem moves establishes the lemma. $\Box$

Let p_k be a k-colored 0-simplex of T, the dual of an n-gem G, and let $j \neq k$ be another color. The p_k *centered j-bisection in T* is the following operation, introduced by Gagliardi in [6]: bisect each 1-simplex e, whose ends are p_k, p', where p' is j-colored, by creating a new 0-simplex p'' in the middle of e; re-color p_k with j, and color the new 0-simplices p'' with k; bisect every n-simplex $S = \{p_k, p', \ldots, w\}$ containing p_k and p', into $S' = \{p_k, p'', \ldots, w\}$ and $S'' = \{p', p'', \ldots, w\}$. This operation clearly produces another colored pseudo-triangulation of the same n-manifold. We denote by $\mathcal{B}(T, p_k, j)$ the result of the p_k-centered j-bisection of the colored pseudo-triangulation T. Observe that, if $Q_k = \{p_k^1, p_k^2, \ldots, p_k^u\}$ is a set of k-colored 0-simplices of T, performing the p_k^i-centered j-bisections in any order produces the same colored pseudo-triangulation denoted by $\mathcal{B}(T, Q_k, j)$: since there are no n-simplices containing two of the p_k^i's, their respective bisections commute.

We need to interpret the bisection operation in dual terms. Let G be an n-gem and $k \neq j$ be distinct colors in N. Also, let $F_{\overline{k}}$ be an $(n-1)$-gem (possibly non-connected) formed by the union of a given non-empty subset of $(N\backslash\{k\})$-residues of G. Trisect each j-edge e of $F_{\overline{k}}$ by creating two new vertices in the interior of e.

If x and y are the vertices of e, denote by x' and y' the new vertices. Moreover, link x to x' by a j-edge, link x' to y' by a k-edge and link y' to y by a j-edge. Define $I = N \backslash \{j, k\}$ and, for each $i \in I$ and each i-edge in $F_{\overline{k}}$ with ends x and z, link x' and z' by an i-edge. The result of these trisections is an n-gem, denoted by $\mathcal{T}(G, F_{\overline{k}}, j)$. It is called the *trisection of the j-edges in an $(N\backslash\{k\})$-subgraph $F_{\overline{k}}$*, or a *trisection* of G. In Fig. 5 we illustrate an important characteristic of the trisection: it can be factored into two phases, namely, the creation of blobs over the j-edges of $F_{\overline{k}}$, producing G°, and the *correction of the I-edges* to go from G° to $\mathcal{T}(G, F_{\overline{k}}, j)$.

Proposition 11 (Bisection × trisection duality). *Let T be a colored pseudo triangulation and G its dual n-gem. Let Q_k be a non-empty subset of the k-colored 0-simplices of T, and let $F_{\overline{k}}$ be the union of the $(N\backslash\{k\})$-residues corresponding to the 0-simplices in Q_k, via duality. Then the dual of $\mathcal{B}(T, Q_k, j)$ is $\mathcal{T}(G, F_{\overline{k}}, j)$.*

Proof. The proof follows from the geometric interpretation of the definitions. We refer to Fig. 6, where $I = N\backslash\{j, k\}$ and $I' = I\backslash\{i\}$. For our argument, i is an arbitrary color in I.

The small black square labelled I' represents a colored $(n-2)$-simplex with color set I'. Each vertex b of $F_{\overline{k}}$, gets a copy, b', in $\mathcal{T}(G, F_{\overline{k}}, j)$, so that b, b' are the ends of a j-edge. Each i-edge of $F_{\overline{k}}$ with ends (b, a) is thickened to a square (i, j)-bigon with vertices (b, b', a', a). Moreover, each j-edge of $F_{\overline{k}}$ with ends (b, c) is trisected by a j-edge (b, b'), a k-edge (b', c') and another j-edge (c', c). Since these facts hold for each $i \in I$, the proof is complete. □

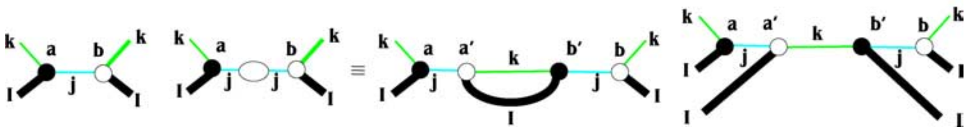

Fig. 5. Factoring trisection by first creating blobs: $G \to G^\circ \to \mathcal{T}(G, F_{\overline{k}}, j)$.

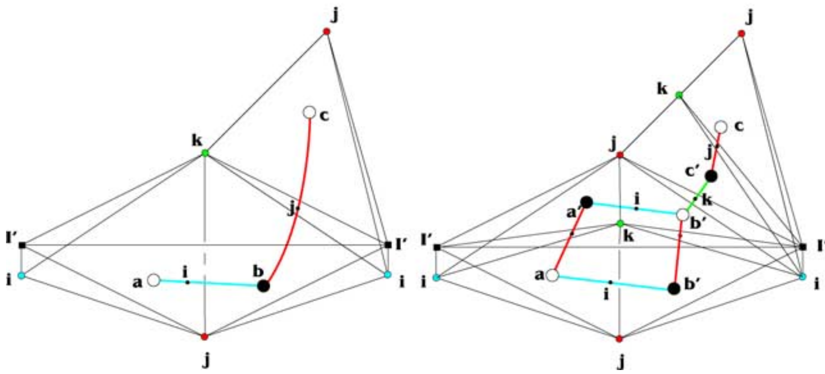

Fig. 6. Bisection/trisection duality.

Casali has proved in [1] that a specific sequence of $(n+1)n/2$ bisections factors the barycentric subdivision of any n-dimensional pseudo-triangulation. The pseudo-triangulation, as well as some subsets of i-colored 0-simplices arising in this factorization, are important to us. So we formalize the method of [1] as an algorithm to provide an adequate terminology, in which the pseudo-triangulations that arise from successive bisections are denoted by T_i^j with $0 \le j \le n$ and $0 \le i \le n-j+1$. The 0-skeleton of a pseudo-triangulation T is denoted by $S_0(T)$. For $0 \le j \le n-1$, $0 \le i \le n-j-1$ the subset $Q_i^j \subset S_0(T_i^j)$ is formed by the i-colored 0-simplices of T_i^j, which are 0-simplices of the original triangulation T. This subset plays an important role in the primal-dual algorithm (Proposition 13). Note that the algorithm depends on the ordering of the colors, which we assume to be the usual one $0, 1, \ldots, n$. Variations in this order produce (non-color) isomorphic $T^\star$'s.

Proposition 12 (Casali's algorithm to factor $T \to T^\star$ by bisections). *The following algorithm produces the barycentric subdivison $T^\star$ of a pseudo triangulation T:*

1. $T_0^0 \leftarrow T$;
2. *for* j *from* 0 *to* $n-1$ *do*
3.   *for* i *from* 0 *to* $n-j-1$ *do*
4. $Q_i^j \leftarrow \{p \in S_0(T) \cap S_0(T_i^j) \mid \mathrm{color}(p) = i+1 \text{ in } T_i^j\}$;
5. $T_{i+1}^j \leftarrow \mathcal{B}(T_i^j, Q_i^j, i)$;
6. $T_0^{j+1} \leftarrow T_{n-j+1}^j$;
7. $T^\star \leftarrow T_0^n$.

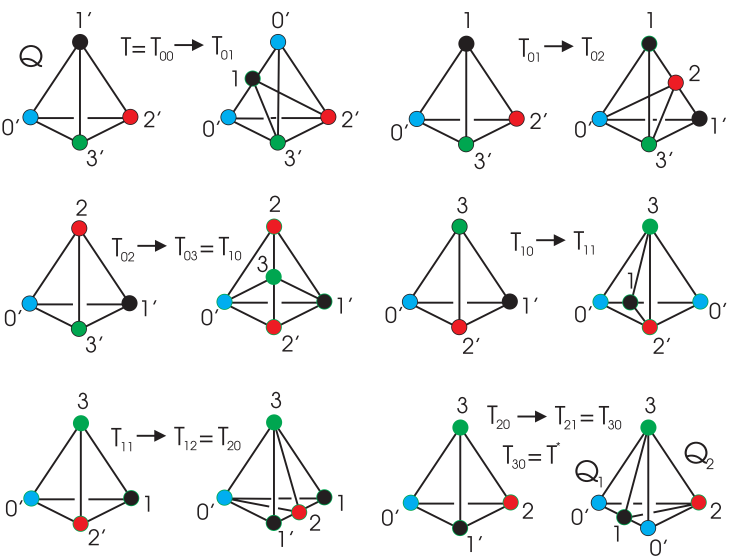

Fig. 7. Snapshot of a 3D-Barycentric Subdivison by $(i, i-1)$-bisections.

Proof. See [1]. □

The final pseudo-triangulation $T^\star$ coincides with the barycentric subdivision of T. In Fig. 7 we present a snapshot of this algorithm in dimension 3 to obtain the barycentric subdivision $T^\star$ of a colored pseudo-triangulation T. The 0-simplices of $T^\star$ are naturally partitioned into $n+1$ classes: $v \in V(T^\star)$ is either an original 0-simplex of T or is in the interior of an original i-simplex of T, $i = 1, 2, \ldots, n$. Such a 0-simplex is said to be *a dimension i representing 0-simplex of $T^\star$*. We mark each original 0-simplex, with a tick in its color label. Each time a new 0-simplex is created, its color and its representing dimension are the same. Only ticked 0-simplices (original 0-simplices) are used as center of a bisection. Each time a 0-simplex of color i is used as center, its color decreases by 1 and it remains ticked. Such a 0-simplex is used i times as center and so its final color is 0. In the figure we display a snapshot of the process for $n = 3$ focusing in a tetrahedron Q, and choosing, after each bisection, one of the two tetrahedra to focus. Of course, the method operates in parallel in all tetrahedra. At the end of the six bisections we arrive display two tetrahedra Q_1 and Q_2 of $T^\star$, the pair being two of the 24 tetrahedra in which Q is subdivided in its barycentric subdivision. The color of each 0-simplex of $T^\star$ coincides with its representing dimension.

To prove our main results, it is essential to use gems instead of pseudo-triangulations. Therefore we rewrite the previous algorithm using gems and their dual colored pseudo-triangulations, emphasizing the use of trisections instead of bisections. The notation $B \overset{d}{\leftarrow} A$ means that B is obtained from A by geometrical duality. The colored pseudo-triangulations T_i^j are carried along in the *primal-dual algorithm*, formalized in Proposition 13, only to produce the $(N\backslash\{i\})$-subgraph $E_{\hat{i}}^j$ dual of Q_i^j. The reason why the algorithm is not entirely in terms of gems is because it is impossible to define adequate $E_{\hat{i}}^j$ without duality.

Proposition 13 (Primal-dual algorithm to factor $F \to F^\star$). *The following algorithm produces the barycentric thickening $F^\star$ of an n-gem F by trisections*:

1. $F_0^0 \leftarrow F$; $T \overset{d}{\leftarrow} F$; $T_0^0 \overset{d}{\leftarrow} F_0^0$;
2. *for j from* 0 *to* $n-1$ *do*
3.   *for i from* 0 *to* $n-j-1$ *do*
4. $Q_i^j \leftarrow \{p \in S_0(T) \cap S_0(T_i^j) \mid \text{color}(p) = i+1 \text{ in } T_i^j\}$;
5. $E_{\hat{i}}^j \overset{d}{\leftarrow} Q_i^j$; *(gem $E_{\hat{i}}^j$ is an $(N\backslash\{i\})$-subgraph of F_i^j)*;
6. $F_{i+1}^j \leftarrow \mathcal{T}(F_i^j, E_{\hat{i}}^j, i)$; $T_{i+1}^j \overset{d}{\leftarrow} F_{i+1}^j$;
7. $F_0^{j+1} \leftarrow F_{n-j+1}^j$; $T_0^{j+1} \overset{d}{\leftarrow} F_0^{j+1}$;
8. $F^\star \leftarrow F_0^n$.

Proof. The proof follows, by duality, from Proposition 12. □

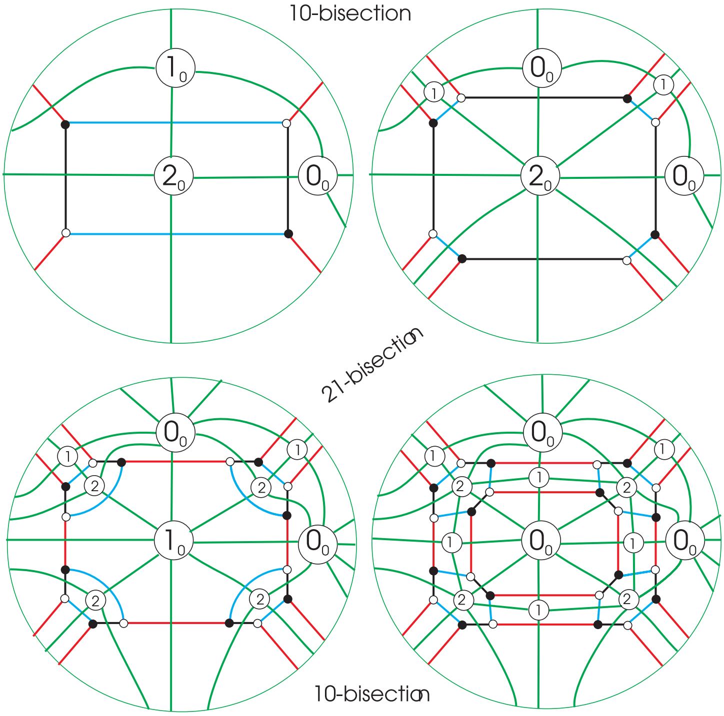

Fig. 8. An example of the primal-dual algorithm to factor the barycentric thickening.

In Fig. 8 we present an application of the primal-dual algorithm applied to the complete graph K_4 embedded into the real projective plane $\mathbb{RP}^2$. This simple example illustrates the fact that the theory applies to non-orientable manifolds as well.

5. Properties of Trisections

As we have mentioned, the operation of trisection $G \mapsto \mathcal{T}(G, F_{\overline{k}}, z)$ can be factored into two phases, namely, creation of blobs in the z-edges of $F_{\overline{k}}$, producing G°, and *correction of the I-edges* to go from G° to $\mathcal{T}(G, F_{\overline{k}}, z)$, replacing j by z (see Fig. 5). Note that this I-correction phase can be performed by a finite number of slimming flips. In this context, see Proposition 16.

After the trisection, each I-residue R contained in $F_{\overline{k}}$ is duplicated and becomes an $(I \cup \{z\})$-prism with bases isomorphic to R in $\mathcal{T}(G, F_{\overline{k}}, z)$. Note that $V(\mathcal{T}(G, F_{\overline{k}}, z)) \supset V(G)$.

Proposition 14 (Correspondence $F_{\overline{k}} \leftrightarrow F'_{\overline{\{z\}}}$). *Let G be an n-gem, $I \cup \{k, z\}$ be a partition of N, $F_{\overline{k}}$ be an $(I \cup \{z\})$-subgraph of G and $H = \mathcal{T}(G, F_{\overline{k}}, z)$. The vertices of H which are not vertices of G* (primed vertices) *are the vertices of an*

$(I\cup\{k\})$-subgraph of H denoted by $F'_{\overline{\{z\}}}$. An $(I\cup\{k\})$-residue R' (a component) of $F'_{\overline{\{z\}}}$ whose vertex set is $\{x'_1, x'_2, \ldots, x'_u\}$, corresponds to an original $(I\cup\{z\})$-residue R (a component) of $F_{\overline{k}}$ with vertex set $\{x_1, x_2, \ldots, x_u\}$. Moreover, by replacing each k-edge of this residue by a z-edge we get an isomorphic copy of R.

Proof. Straightforward consequence of the definition of $\mathcal{T}(G, F_{\overline{k}}, z)$. □

For gems with a color partition $I\cup J\cup\{k,z\}$, the operation of trisection interacts with dipoles in a way that merits further study. We need the following proposition in the proof of the main lemma, Proposition 17.

Proposition 15 (Dipoles and trisection). *Let G be an n-gem, $I\cup J\cup\{k,z\}$ be a partition of N, $F_{\overline{k}}$ be an $(I\cup J\cup\{z\})$-subgraph of G and $H = \mathcal{T}(G, F_{\overline{k}}, z)$.* (1) *if $a^G_{I\cup\{k\}}$ is a dipole, then $a^H_{I\cup\{k\}}$ is a dipole*; (2) *if $a^G_{I\cup\{k,z\}}$ is a dipole, then $a^H_{I\cup\{k\}}$ and $(a')^H_{I\cup\{k\}}$ are dipoles.*

Proof. Under the hypothesis of (1), assume that $a^H_{I\cup\{k\}}$ is not a dipole. Take a minimal path π, linking a to b in H and having only $(J\cup\{z\})$-edges. We claim that π uses z-edges, otherwise π links a to b in H having only J-edges. Such a path is also a path in G, which contradicts the hypothesis of (1). So the concatenated path $\pi = \pi_1 \circ u_{\{z\}} \circ \pi'_2 \circ w_{\{z\}} \circ \pi_3$, with π_1 disjoint from $F'_{\overline{\{z\}}}$ and π'_2 links u' to w' in a component R' of $F'_{\overline{\{z\}}}$ by J-edges. A J-path in R linking u to w corresponds to π'_2. It follows that $\pi_1 \circ \pi_2 \circ \pi_3$ links a to b in H by $(J\cup\{z\})$-edges, contradicting the minimality of π. The proof of (1) is complete.

Under the hypothesis of (2), assume that $a^H_{I\cup\{k\}}$ is not a dipole. Take a minimal $(J\cup\{z\})$-path π linking a to b in H. We claim that π uses z-edges, otherwise π is a J-path linking a to b in H. Such a path is also a path in G, which contradicts the fact that $a^G_{I\cup\{k,z\}}$ is a dipole. So, $\pi = \pi_1 \circ u_{\{z\}} \circ \pi'_2 \circ w_{\{z\}} \circ \pi_3$, with π_1 disjoint

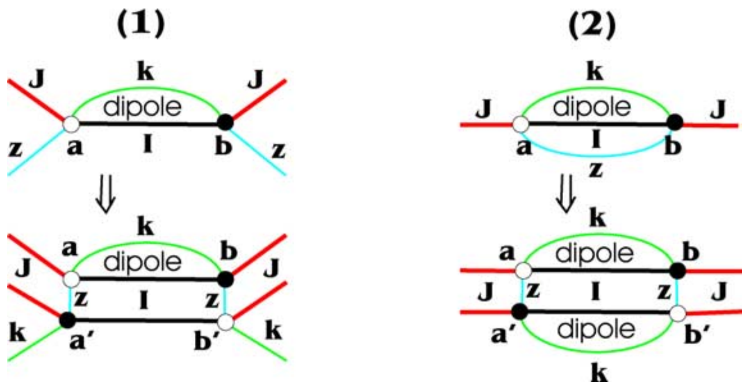

Fig. 9. Dipoles and trisection.

from $F'_{\overline{\{z\}}}$ and π'_2 links u' to w' in a component R' of $F'_{\overline{\{z\}}}$ by J-edges. A J-path in R linking u to w corresponds to π'_2. It follows that $\pi_1 \circ \pi_2 \circ \pi_3$ links a to b in H by $(J \cup \{z\})$-edges, contradicting the minimality of π. Then $a^H_{I\cup\{k\}}$ is a dipole. If π' is a $(J \cup \{z\})$-path which links a' to b' in H, then $a_{\{z\}} \circ \pi' \circ b_{\{z\}}$ is a $(J \cup \{z\})$-path in H linking a to b. Since $a^H_{I\cup\{k\}}$ is a dipole, such π' does not exist and $(a')^H_{I\cup\{k\}}$ is a dipole, proving (2). □

If we do not care about maintaining the induced manifold while going from G to $\mathcal{T}(G, F_{\overline{k}}, z)$, then it is easy to accomplish the passage. A *fake* s-flip is the slimming of a dipole which produces a non-dipole.

Proposition 16 (Blobs and fake s-flips). *The passage $G \to \mathcal{T}(G, F_{\overline{k}}, z)$ can be accomplished by means of blob creations followed by (possibly fake) s-flips.*

Proof. Start by creating a blob at each z-edge of $F_{\overline{k}}$, producing an n-gem $G^\circ = G_0$. Label the new vertices with a prime, such that v' is linked to v by a z-edge. Let $\ell \leftarrow 0$ and iterate the following procedure. Consider in G_ℓ a vertex v' in one of the created blobs and a color c, where $c \in I = N\backslash\{z, k\}$. Suppose that the c-neighbor of v' is u' and that v has, as c-neighbor, the vertex w. If $u \neq w$, we say that v' is *c-wrong* in G_ℓ. Otherwise, we say that v' is *c-correct*. Assume that v' is c-wrong. Let x' be the c-neighbor of w'. The *(v', c)-correcting flip* replaces the c-edges (u', v') and (w', x') by the c-edges (v', w') and (u', x'), and defines $G_{\ell+1}$. Note that v' and w' are both c-wrong in G_ℓ and v' is c-correct in $G_{\ell+1}$. Thus the total number of c-wrong vertices is smaller in $G_{\ell+1}$ than in G_ℓ. Make $\ell \leftarrow \ell + 1$ and repeat until there are no c-wrong vertices. The final G_ℓ is clearly, by definition, $\mathcal{T}(G, F_{\overline{k}}, z)$. □

If, in the above proof, at the (v', c)-correcting flip, we could make sure that the set of edges between u' and v' forms a dipole, then this flip would be a genuine dipole slimming flip (an s-flip), and we would have completed the factorization $G \to \mathcal{T}(G, F_{\overline{k}}, z)$ by gem moves. However, that set of edges is not in general a dipole, and we just have a *fake s-flip*. A serious problem with the construction of the proof of Proposition 16 is because a fake s-flip, applied to a gem, produces in general a pseudogem (or possibly a gem of a different manifold) and so it is impossible to factor these flips by gem moves. We have seen in Proposition 3, however, that a clean flip factors into a pair of dipole moves. In view of Proposition 3, to factor $G \to \mathcal{T}(G, F_{\overline{k}}, z)$ by gem moves it would be sufficient to perform the construction of Proposition 16 using only clean flips. However, this is impossible for arbitrary gems G's and $(N\backslash\{k\})$-subgraphs $F_{\overline{k}}$'s. In the rest of this paper we show that, in a very particular class of gems and $(N\backslash\{k\})$-subgraphs, gem moves are sufficient to perform the correcting flips. Moreover, this sufficiency, in this restricted class, implies that $G \to \mathcal{T}(G, F_{\overline{k}}, z)$ is factorable by gem moves in general.

In dimension 2 the above algorithm, with an appropriate ordering of the corrections, produces only true s-flips. This could be used to establish the basis of the inductive proof of the main lemma. We do not use it because we will take the case $n = 1$ as the basis. There are examples showing that, for dimensions $n \geq 3$, no ordering of corrections is adequate. However, in the final section we establish the existence of such an adequate ordering for an arbitrary dimension n, assuming the hypothesis that $F_{\overline{k}}$ is itself the barycentric thickening of some $(n-1)$-gem. Finding this hypothesis was the key step in achieving the proof of the main lemma.

6. The Main Lemma

A dipole D of some n-gem is said to be *associated to a gem move* if it is the blob (an n-dipole) which is being cancelled or created, or if it is the dipole which is being thickened or slimmed. Consider a single gem move μ, yielding n-gem G'' from gem G', having D as its associated dipole. Assume that G is a gem satisfying $V(G) \subset V(G') \cap V(G'')$. If $V(D) \cap V(G) = \emptyset$, then μ is a *G-special gem move.*

Proposition 17 (The main lemma). *Let G be an n-gem and $H = \mathcal{T}(G, F_{\overline{k}}, z)$, where $F_{\overline{k}}$ is an $(N\backslash\{k\})$-residue of G. Then there is a sequence of G-special gem moves linking G to H.*

Proof. The proof is by induction on the dimension n of the manifold induced by G. For $n = 1$, the trisection is accomplished by blob creations on the z-edges. No correction phase is needed.

To simplify the notation in proving the induction step, let us relabel the colors, so that $k = n$ and $z = n - 1$. Also denote $F_{\overline{k}}$ by F. However, in our next use of the inductive hypothesis, we are allowed to assume that the proposition is true for the $(n-1)$-case with arbitrary colors k and z.

By using the factorization of Proposition 13 and the inductive hypothesis, there exists a sequence of gems from F to its barycentric thickening $F^\star$, namely, $\sigma_F = (F = F_0, F_1, F_2, \ldots, F_p = F^\star)$ so that F_i is obtained from F_{i-1} by a single F-special gem move. Note that σ_F is obtained by refining the sequence (F_i^j) defined by the algorithm of Proposition 13. This refinement is allowed by the inductive hypothesis. Starting with $G_0 = G$ we construct a sequence $\sigma_G = (G = G_0, G_1, G_2, \ldots, G_p = \overline{G})$: each time that we create or cancel a blob in some F_ℓ we create or cancel a blob in G_ℓ with an extra n-edge. Since the moves in σ_F are F-special, a dipole slimming or thickening in F_ℓ is also a dipole (the same dipole with an extra n-edge) slimming or thickening in G_ℓ. So σ_G is well defined and it follows that we can go from G to $\overline{G}$ by G-special gem moves. Observe that F_ℓ is an $(N\backslash\{n\})$-residue of G_ℓ. Define $H_\ell = \mathcal{T}(G_\ell, F_\ell, n-1)$, $0 \leq \ell \leq p$, and consider $\sigma_H = (H = H_0, H_1, H_2, \ldots, H_p = \overline{H})$. We claim that we can go from $H_{\ell-1}$ to H_ℓ by a sequence of G-special gem moves. We also claim that we can go from $\overline{G}$ to $\overline{H}$ by means of a finite sequence of gem moves, which do not involve the vertices of G. Up to these two claims, the proof of the main lemma is established. We prove the first of these claims here. The proof of the second one is given in the last section of the paper.

First suppose that G_ℓ is obtained from $G_{\ell-1}$ by a clean flip. Let $\{G', G''\} = \{G_{\ell-1}, G_\ell\}$ so that G'' is obtained from G' by thickening a dipole D. As before, D contains an n-edge. We have three cases:

(1) The color h of the edges which are being flipped is $n-1$;
(2) $h \neq n-1$ and D has an $(n-1)$-edge;
(3) $h \neq n-1$ and D has no $(n-1)$-edge.

Case 1 is dealt with in Fig. 10. Step $H' \xrightarrow{\text{Fig.10}} H'_1$: since $a^{G'}_{\{n\}\cup I}$ is a dipole, so is $a^{H'}_{\{n\}\cup I}$ by part (1) of Proposition 15. Thicken the latter to obtain dipole $a^{H'_1}_{K\cup I\cup\{n-1\}}$.

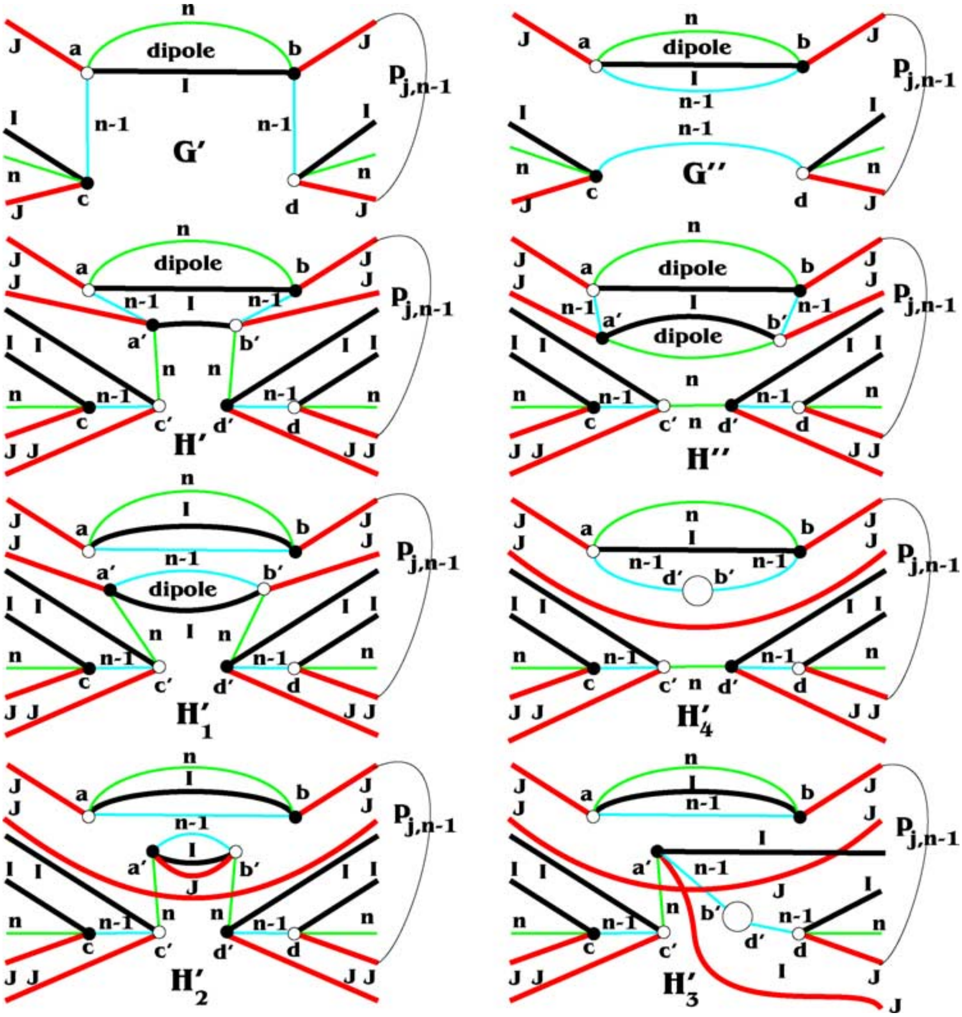

Fig. 10. Linking $(z, F_{\overline{k}})$-trisections of G', G'' differing by a clean flip: case 1.

Step $H'_1 \xrightarrow{\text{Fig.10}} H'_2$: Proposition 14 implies that $(a')^{H'_1}_{I\cup\{n-1\}}$ is a dipole: $(a')^{H'_1}_{\{n\}\cup J}$ and $(b')^{H'_1}_{\{n\}\cup J}$ are distinct because $a^{G'}_{\{n-1\}\cup J}$ and $b^{G'}_{\{n-1\}\cup J}$ are too. Perform J-thickening of the dipole $(a')^{H'_1}_{I\cup\{n-1\}}$. Step $H'_2 \xrightarrow{\text{Fig.10}} H'_3$: note that $(b')^{H'_2}_{\{n\}}$ is a dipole. Do the $(I \cup J)$-thickening of this dipole, producing an $(n-1)$-blob at H'_3. Alternatively, the same result can be obtained by rotating the n-blob $(b')^{H'_2}_{\{n\}}$ to the $(n-1)$-edge incident to d'. Step $H'_3 \xrightarrow{\text{Fig.10}} H'_4$: transfer the blob along the path $p_{j,n-1}$. Note that $J \neq \emptyset$, otherwise G'' would have one more component than G', which is not possible via a gem move. The path $p_{j,n-1}$ is a path linking d to b, which alternates $(n-1)$-edges and j-edges, for a certain $j \in J$. Step $H'' \xrightarrow{\text{Fig.10}} H'_4$: by part (2) of Proposition 15, $(a')^{H''}_{\{n\}\cup I}$ is a dipole, because $a^{G''}_{I\cup\{n-1,n\}}$ is a dipole. The J-thickening of $(a')^{H''}_{\{n\}\cup I}$ yields H'_4. Note that, except maybe for c and d, all the vertices appearing in Fig. 10 are not in G. Therefore, all the gem moves performed to go from H' to H'' are G-special. In particular, the moves which rotate (many times) the blob with vertices b' and d' are G-special. This establishes case 1.

Case 2 is dealt with in Fig. 11. Step $H' \xrightarrow{\text{Fig.11}} H'_1$: By Proposition 15, part (2) with the J of that proposition replaced by the current $J \cup \{h\}$, we obtain that $(a')^{H'}_{I\cup\{n\}}$ is a dipole. Thicken the latter by $\{h\}$, obtaining the dipole $a^{H'_1}_{K\cup I\cup\{h\}}$. Step $H'_1 \xrightarrow{\text{Fig.11}} H'_2$: apply J-thickening in the latter dipole, yielding an $(n-1)$-blob. Step $H'_2 \xrightarrow{\text{Fig.11}} H'_3$: apply two rotations to this blob. Step $H'_3 \xrightarrow{\text{Fig.11}} H'_4$: Note that $a^{H'_3}_{\{h\}\cup J} = a^{G'}_{\{h\}\cup J} \neq b^{G'}_{\{h\}\cup J} = b^{H'_3}_{\{h\}\cup J}$. Thus $a^{H'_3}_{I\cup K\cup\{n-1\}}$ is a dipole. Perform the $\{h\}$-thickening of this dipole, followed by the transfer of the $(n-1)$-blob at d along the p_{hj} to the $(n-1)$-edge at a. Note that $J \neq \emptyset$, otherwise G'' would have one more component than G', which is not possible via a gem move. The path p_{hj} is a path linking d to b, which alternates j-edges and h-edges, for a certain $j \in J$. Step $H'' \xrightarrow{\text{Fig.11}} H'_4$: by part (2) of Proposition 15 with $I \cup \{h\}$ in place of I, $(a')^{H''}_{I\cup\{n,h\}}$ is a dipole, because $a^{G''}_{I\cup\{h,n-1,n\}}$ is a dipole. Do the J-thickening of this dipole to get H'_4. Note that, as in the previous case, except maybe for c and d, all the vertices appearing in Fig. 11 are not in G. Therefore, all the gem moves performed to go from H' to H'' are G-special. In particular the moves which rotates (many times) the blob with vertices a' and b' are G-special. The proof of case 2 is complete.

For the proof of case 3 we use Fig. 12. Step $H' \xrightarrow{\text{Fig.12}} H'_1$: by Proposition 15, part (1) with J of that proposition replaced by the current $J \cup \{h\}$, we obtain that $a^{H'}_{I\cup\{n\}}$ is a dipole. Do the $\{h, n-1\}$-thickening of the latter to obtain dipole $a^{H'_1}_{I\cup\{h,n-1,n\}}$. Step $H'_1 \xrightarrow{\text{Fig.12}} H'_2$: by Proposition 14 $(a')^{H'_1}_{I\cup\{n-1\}}$ is a dipole. Do the $\{h\}$-thickening of this dipole. Step $H'' \xrightarrow{\text{Fig.12}} H'_2$: by part (1) of Proposition 15, with the current $I \cup \{h\}$ in place of I, $a^{H''}_{I\cup\{h,n\}}$ is a dipole. The $\{n-1\}$-thickening of this dipole yields H'_2. Observe that the only vertices which might be in G are c and d. Therefore, all the gem moves performed from H' to H'' are G-special, establishing case 3.

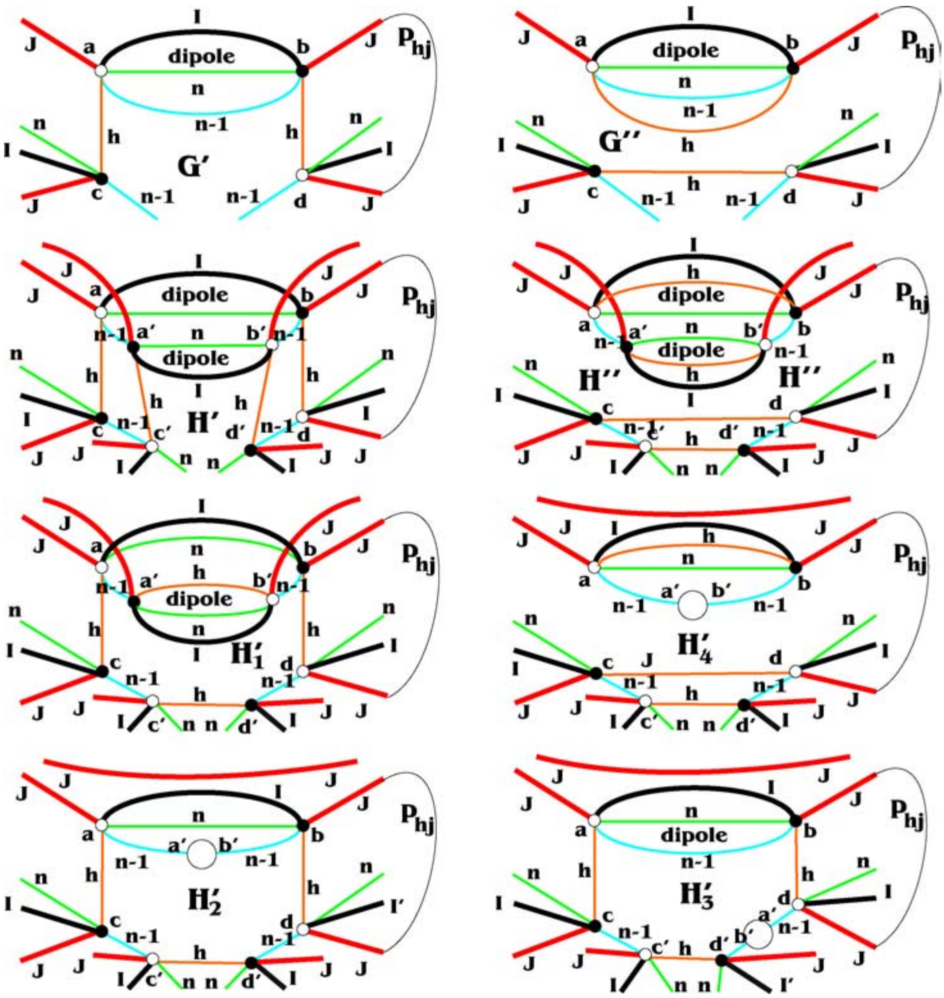

Fig. 11. Linking $(n-1, F_{\overline{k}})$-trisections of G', G'' differing by a clean flip: case 2.

Suppose now that G_ℓ is obtained from $G_{\ell-1}$ by a G-special blob cancellation or creation, that $\{G', G''\} = \{G_{\ell-1}, G_\ell\}$ and that G' contains the blob D to be cancelled. Assume first that D is over an h-edge, $h \neq n-1$. By the definition of G' from F', D contains an n-edge and so $h \notin \{n-1, n\}$. We refer to Fig. 13. To go from H' to H'' we have thickened the dipole $e^{H'}_{I \cup \{n\}}$ using the $(n-1)$-colored edge. Since e, f are not vertices in $V(G)$ it follows that this thickening is G-special. After that, two h-blobs arise and their cancelling produces H''. These moves are G-special, because $\{e, f, e', f'\} \cap V(G) = \emptyset$. Next, consider the case in which the blob D is over an $(n-1)$-edge. The creation of this $(n-1)$-blob can be factored as the creation of an h-blob, followed by a pair (s-flip,t-flip) performing the rotation of the h-blob into an $(n-1)$-blob, as explained in Proposition 4. Each one of these

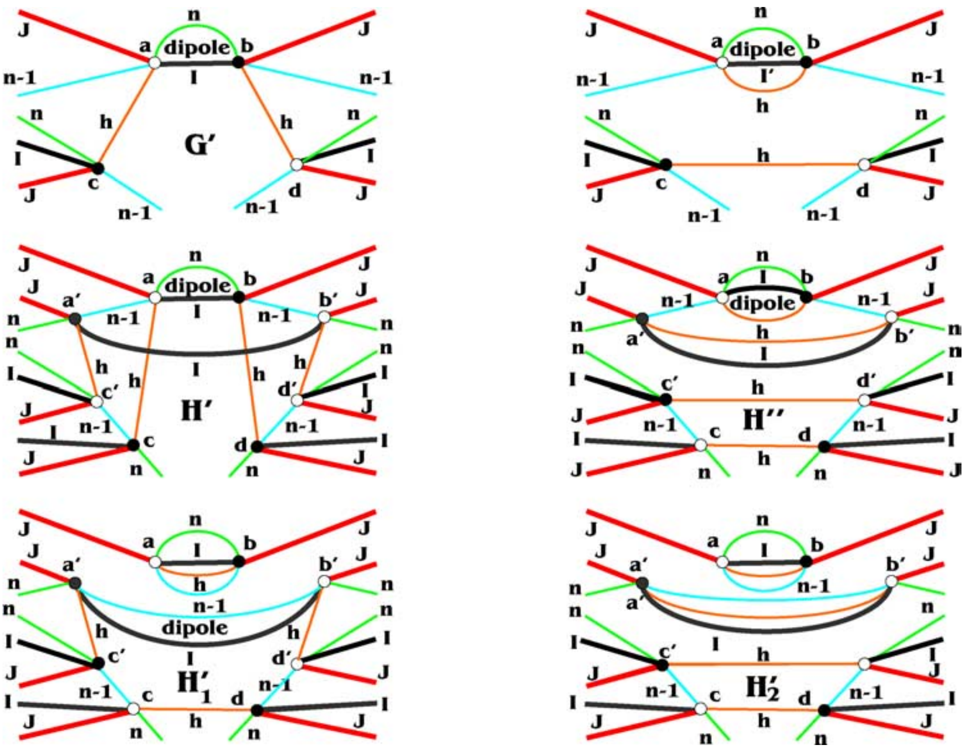

Fig. 12. Linking $(n-1, F_{\overline{k}})$-trisections of G', G'' differing by a clean flip: case 3.

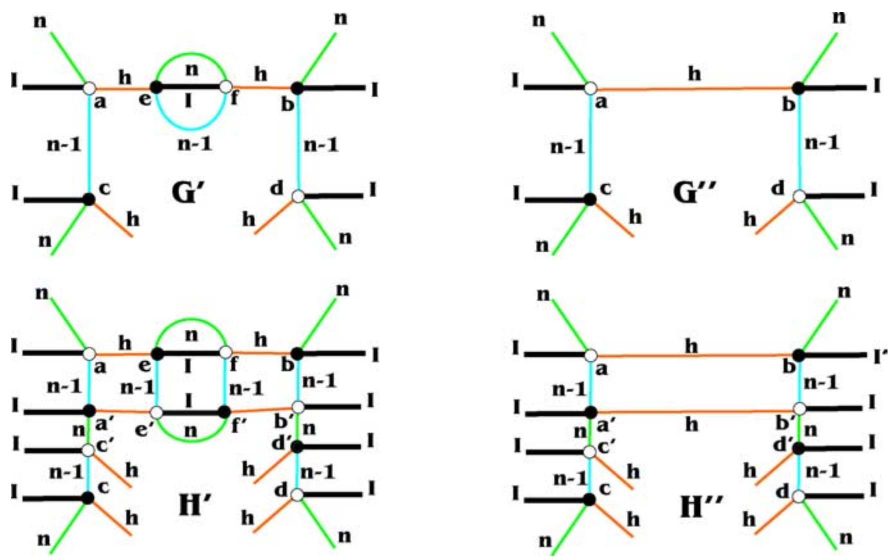

Fig. 13. Linking, by gem moves, $(z, F_{\overline{k}})$-trisections of gems differing by an h-blob, $h \neq n$.

3 special G-moves were treated in previous cases. Thus, we can go from $\overline{G}$ to $\overline{H}$ by a finite sequence of G-special gem moves.

To finish the proof of the main lemma we still need to establish that $\overline{G}$ is linked to $\overline{H}$ by a finite sequence of G-special gem moves. □

We write $G' \ (\uparrow \square \downarrow)^*_G \ G''$, if $V(G') \cap V(G'') \supseteq V(G)$ and G'' is obtained from G' by a finite sequence of G-special gem moves. In Sec. 8 we conclude the proof of the main lemma by establishing that $\overline{G} \ (\uparrow \square \downarrow)^*_G \ \overline{H}$. In the next section we prepare the ground for this final step.

7. The Permutohedron

The tools to define the ingredients used in the proof that $\overline{G} \ (\uparrow \square \downarrow)^*_G \overline{H}$ are the properties of a convex $(n-1)$-polytope P^{n-1} embedded into $\mathbb{R}^{n-1}$. For a positive integer q, the *q-permutohedron*, see [2], is a $(q-1)$-gem with color set $\{0, 1, \ldots, q-1\}$ inducing the sphere S^{q-1}, defined as follows: its vertex set is the set of the $(q+1)!$ permutations of $\{0, 1, \ldots, q\}$. Given two vertices π, and π' there is an $(q-1-i)$-colored edge linking them if π' is obtained from π by interchanging its ith and $(i+1)$th symbols, $i = 0, 1, \ldots, q-1$. This accounts for all the $\frac{q(q+1)!}{2}$ edges of P^q, concluding its definition. A simple structural property of the permutohedron follows:

Proposition 18 (Bi-colored polygons of size 4 and 6 in the permutohedron). *For $q \geq 1$ and $0 \leq i < j \leq q$, each $\{i, j\}$-residue in P^q has either 4 or 6 vertices according to $j - i > 1$ or $j - i = 1$.*

Proof. From the definition of the i- and j-edges of P^{n-2}, if $j > i+1$, the transpositions $(i, i+1)$ and $(j, j+1)$ commute, forming an orbit with 4 elements. If $j = i+1$, the orbit has 6 elements. □

Our next target is to define an ε-canonical embedding of the q-permutohedron, for $q \geq 1$, as a subset $P^q_\varepsilon \subseteq \mathbb{R}^q$, so that each residue of P^q embeds as the 1-skeleton of a convex polytope contained in P^q_ε.

First we need to embed the regular q-simplex into $\mathbb{R}^q$. Denote by T^q the regular q-simplex whose 1-faces have length 1 embedded in $\mathbb{R}^q$. We want to assign fixed well defined *canonical coordinates* for the 0-faces (the vertices) T^q. Essentially, to make computations easier, in the canonical coordinates, we make the barycenter of T^q coincide with the origin of $\mathbb{R}^q$. The construction is by induction. The coordinates of the vertices $u_{1,0}, u_{1,1}$ of T^1 in $\mathbb{R}^1$ are respectively $-1/2$ and $1/2$. Having defined the canonical coordinates of the vertices $u_{q-1,0}, u_{q-1,1}, \ldots, u_{q-1,q-1}$ of T^{q-1} in $\mathbb{R}^{q-1}$, the coordinates of the vertices of T^q in $\mathbb{R}^q$ are defined as follows. Denote by h_q the *height of the regular q-simplex*, which is $h_q = \sqrt{1 - \|u_{q-1,0}\|^2}$. Let $u'_{q,q}$ of $\mathbb{R}^q$ be $(0, 0, \ldots, 0, h_n)$. This point, together with the vertices of T^{q-1}, to whose coordinates we append an extra 0, define a regular q-simplex $(T')^q$. To make it canonical we just need to apply the translation by $t_q = (0, 0, \ldots, 0, -h_q/q)$. Indeed, the q-th coordinate of the barycenter of $(T')^q$ is h_q/q. After the translation on $(T)'^q$ by t_q, which defines T^q, the barycenter of the latter is the origin of $\mathbb{R}^q$, as desired. In this way, define $u_{q,j} = (u_{q-1,j}, 0) + t_q$, $1 \leq j \leq q-1$, and $u_{q,q} = u'_{q,q} + t_q$. The definition of the canonical coordinates of T^q is concluded. These coordinates are used to embed the permutohedron P^q into $\mathbb{R}^q$, as follows.

The parameter ε is a real number $0 < \varepsilon \leq 1/q$. The definition is inductive. $P^1_\varepsilon = T^1$. Suppose $P^{q-1}_\varepsilon \subseteq \mathbb{R}^{q-1}$ is defined. Let $P^{q-1}_{\varepsilon,0} \subseteq \mathbb{R}^q$ be constructed from $P^{q-1}_\varepsilon \subseteq \mathbb{R}^{q-1}$ as follows. Append an extra 0 to the coordinates of the vertices of P^{q-1}_ε to consider a subset of $\mathbb{R}^q$. Let $M_q = \varepsilon \cdot P^{q-1}_\varepsilon + u_{q,q}$, where $u_{q,q}$ is the qth vertex of T^q. Thus, M_q is an ε-scaled and $u_{q,q}$-translated copy of P^{q-1}_ε. The label $\pi \in S_q$ of a vertex of P^{q-1}_ε corresponds to the label $\pi' \in S_{q+1}$ of the corresponding vertex of M_q, obtained from π by prefixing q to it. From M_q we obtain M_j, $0 \leq j \leq q-1$, as its image under the reflection $\rho_{q,j}$ along the hyperplane orthogonal to the line linking the barycenter of T^q to the barycenter of its $(q-2)$-dimensional face, induced by all the vertices of T^q except $u_{q,q}$ and $u_{q,j}$. Note that reflection $\rho_{q,j}$ interchanges these two vertices. We label the vertex of M_j corresponding to a vertex π of M_q by interchanging the symbols q and j. In consequence, each vertex π of M_j satisfies $\pi(0) = j$. All vertices of P^q have been positioned in $\mathbb{R}^q$. The *ε-canonical embedding* $P^q_\varepsilon \subseteq \mathbb{R}^q$ of P^q is

$$P^q_\varepsilon = \text{conv.hull}\left(\bigcup M_j \mid 0 \leq j \leq q\right), \tag{1}$$

where conv.hull(S) denotes the convex hull of a set $S \subseteq \mathbb{R}^q$ [17]. See Fig. 14 for an illustration of the case $q = 2$ of this construction. We emphasize that P^q denotes an abstract (not embedded) $(q-1)$-gem while P^q_ε denotes the same gem ε-canonically embedded into $\mathbb{R}^q$. The embedding depends only on the parameter $\varepsilon \in (0, 1/q]$. Note

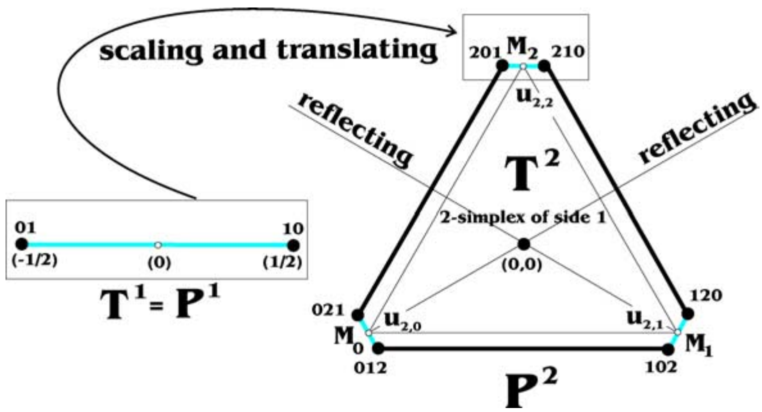

Fig. 14. Construction of P^2_ε from T^2 and P^1: canonical embedding for P^2 in $\mathbb{R}^2$.

that $\lim_{\varepsilon\to 0} P^q_\varepsilon = T^q$. Moreover, from the symmetry of the recursive construction, any two J-residues of P^q_ε are isometric, for arbitrary $J \subseteq \{0,1,\ldots,q\}$.

Given a q-gem F, we show how to obtain its barycentric thickening $F^\star$ combinatorially from F and $|V(F)|$ q-permutohedra, $q \geq 1$. Consider a collection of disjoint q-permutohedra in 1-1 correspondence with the vertices of F, so that a vertex v corresponds to P^q_v (here the subscript v has nothing to do with ε in the embedding P^q_ε). The collection of permutohedra accounts for the $(\{0,1,\ldots,q-1\})$-subgraph of a q-gem $F^\diamond$. To conclude the definition of $F^\diamond$ it is enough to define the q-colored edges. If v and w are linked in F by an i-edge, $i \in \{0,1,\ldots,q\}$, for each $\{0,1,\ldots,q\}$-permutation π such that $\pi(0) = i$, link the vertex π of P^q_v to the vertex π of P^q_w by a q-colored edge $[(v,\pi),(w,\pi)]$. The *secondary color* of this q-edge is i, the original color of (v,w) in F. Let $F^\diamond$ be the q-gem obtained by considering the collection of q-permutohedra linked by the q-edges.

An illustration of this construction is given in Fig. 15, for the case where F is a 2-gem, the 3-edge colored 1-skeleton of a cube. As a matter of fact, in general, $F^\diamond$ coincides with the barycentric thickening of F, as we now prove.

Proposition 19 (Equality $F^\diamond = F^\star$). *Consider the q-gem $F^\diamond$, $q \geq 1$, just defined. Then $F^\diamond = F^\star$.*

Proof. The combinatorial description of the barycentric thickening $F^\star$ comes from duality and is the following. The vertices of $F^\star$ are the sequences (v,π), where v is a vertex of F and π is a permutation of the color set $\{0,1,\ldots,q\}$. We interpret (v,π) as the ascending chain of residues: $(v_\emptyset, v_{\{\pi(0)\}}, v_{\{\pi(0),\pi(1)\}}, \ldots, v_{\{\pi(0),\pi(1)\ldots,\pi(q-1)\}})$. Two vertices represented by these sequences are linked by an edge of color $(q-i)$ if the sequences differ only in the i-th coordinate, $i = 0,1,\ldots,q$. This completes the definition of the q-gem $F^\star$. In this way, there is a 1-1 correspondence between the vertices of $F^\star$ and of $F^\diamond$.

Suppose that v is linked to w by a q-edge in F. The diagram below shows that the 1-1 correspondence between the vertex sets preserves color q. This is because the two sequences on the right differ only in their first entries, namely, $v_\emptyset \neq w_\emptyset$. All the subsequent entries are equal, because $v_{\pi(0)} = w_{\pi(0)}$.

$$\begin{array}{ccc}
(v,\pi(0),\pi(1),\ldots,\pi(q)) & \leftrightarrow & (v_\emptyset, v_{\{\pi(0)\}}, v_{\{\pi(0)\pi(1)\}}, \ldots, v_{\{\pi(0)\pi(1)\ldots\pi(q-1)\}}) \\
\updownarrow \textit{ color } q & & \updownarrow \textit{ color } q \\
(w,\pi(0),\pi(1),\ldots,\pi(q)) & \leftrightarrow & (w_\emptyset, w_{\{\pi(0)\}}, w_{\{\pi(0)\pi(1)\}}, \ldots, w_{\{\pi(0)\pi(1)\ldots\pi(q-1)\}})
\end{array}$$

To prove that color $q-i$, $1 \leq i \leq q$, is also preserved by the correspondence, consider the diagram below.

$$\begin{array}{ccc}
(v,\ldots,\pi(i),\pi(i+1),\ldots) & \leftrightarrow & (v_\emptyset,\ldots,v_{\{\pi(0)\ldots\pi(i)\}}, v_{\{\pi(0)\ldots\pi(i)\pi(i+1)\}},\ldots) \\
\updownarrow \textit{ color } q-i & & \updownarrow \textit{ color } q-i \\
(v,\ldots,\pi(i+1),\pi(i),\ldots) & \leftrightarrow & (v_\emptyset,\ldots,v_{\{\pi(0)\ldots\pi(i+1)\}}, v_{\{\pi(0)\ldots\pi(i+1)\pi(i)\}},\ldots)
\end{array}$$

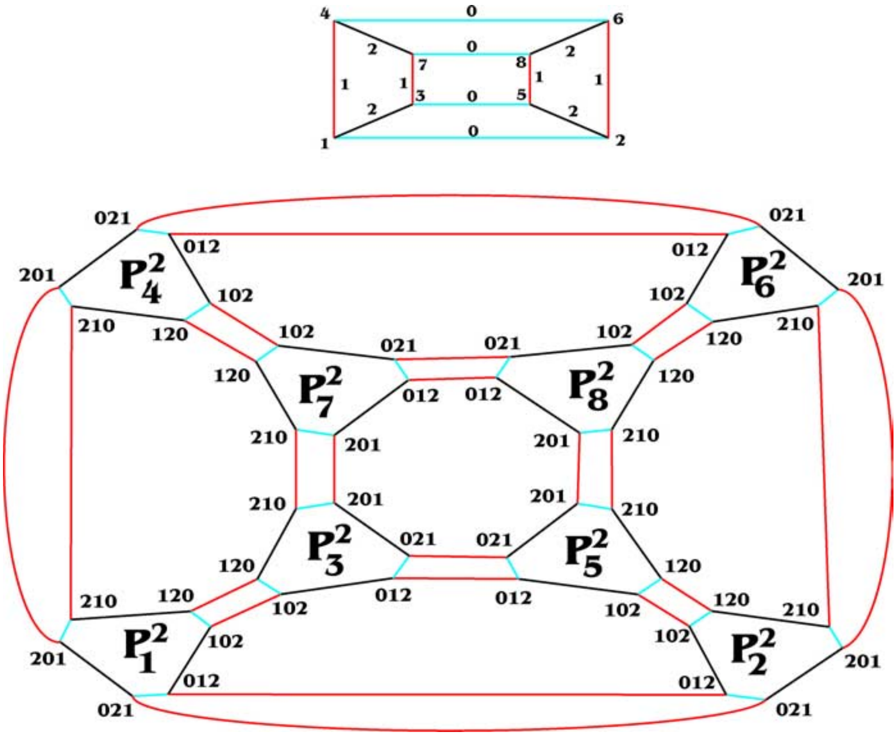

Fig. 15. The construction of $F^\diamond = F^\star$ from F and 2-permutohedra.

Just note that the interchange of $\pi(i)$ and $\pi(i+1)$ on the left sequences corresponds to a difference on the i-entry, $1 \le i \le q$, of the sequences on the right. Thus, color $q-i$ is preserved under the 1-1 correspondence of the vertices of $F^\diamond$ and $F^\star$. So we can conclude that $F^\diamond = F^\star$. □

Let $J = \{j_0, j_1, \ldots, j_\ell\}$ with $j_0 < j_1 < \cdots < j_\ell$. Define $C_\pi^{\{j_0\}}$ as the edge (π, π_{j_0}) of P_ε^q. Inductively define, for $\ell \ge 1$,

$$C_\pi^J = \text{conv.hull}\left(C_\pi^{J\setminus\{j_\ell\}} \cup C_{\pi_{j_\ell}}^{J\setminus\{j_\ell\}} \cup \left(\bigcup_{0\le i\le \ell-1} C_{\pi_{j_i j_\ell}}^{J\setminus\{j_\ell\}}\right)\right), \tag{2}$$

where π_j means the j-neighbor of the vertex π in P_ε^q, and $\pi_{j_i j_\ell}$ means $(\pi_{j_i})_{j_\ell}$. We also define $C_\pi^\emptyset$ as the vertex π of P_ε^q. This extreme situation accounts for the case $\ell = -1$ in the following proposition.

Proposition 20 (Residues R_π^J embedded as convex subpolytopes $\subseteq P_\varepsilon^q$). *Let $J \subseteq \{0, 1, \ldots, q\}$ and $\pi \in V(P_\varepsilon^q)$. Then $C_\pi^J \subseteq P_\varepsilon^q$ is an $(\ell+1)$-dimensional convex polytope corresponding to the J-residue R_π^J of P_ε^q which contains the vertex π.*

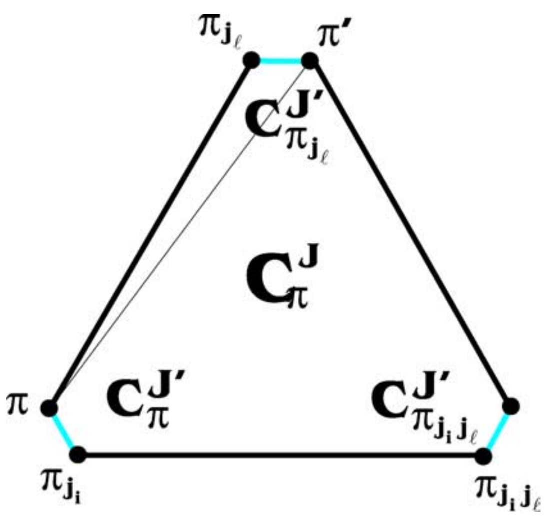

Fig. 16. The convex polytope C^J_π associated to R^J_π, the J-residue of P^q_ε containing π.

This correspondence, in particular, implies that R^J_π *and* C^J_π *have the same vertices and the same edges.*

Proof. The proof is by induction on $|J|$. The result is clear if $|J| \leq 1$. Assume it is true for J' with $|J'| < |J|$. From (2), it is enough to show that the j_ℓ-edge $(\gamma, \gamma_{j_\ell})$ with $\gamma \in V(C^J_\pi)$ is in the boundary of C^J_π, whereas every line from γ to $\gamma' \neq \gamma_{j_\ell}$, with γ and γ' vertices in distinct convex sets of the $(\ell+2)$-convex set union in (2), has each one of its interior points in the interior of C^J_π. This facts follows from the geometry of the definitions. In Fig. 16 we depict a cross section of the situation for $\gamma = \pi$. The case of a generic γ is equivalent up to symmetry. □

Let P_0 be a ε-scaled copy of P^{n-1}_ε, that is, $P_0 = \varepsilon \cdot P^{n-1}_\varepsilon$ and let P_ℓ, $1 \leq \ell \leq |V(F)|$, be a copy of P_0 by a rigid motion in $\mathbb{R}^{n-1}$. Also, let $\mathcal{H}_\ell$, with $1 \leq \ell \leq |V(F)|$, be a hyperplane in $\mathbb{R}^{n-1}$, so that $P_\ell \cap \mathcal{H}_\ell = \emptyset$ and the distance to $\mathcal{H}_\ell$ of the vertices of P_ℓ are all distinct; we say that $\mathcal{H}_\ell$ is *distinctive* for P_ℓ. In this case $\mathcal{H}_\ell$ induces an orientation for the edges of P_ℓ, from its *tail*, the vertex at the longest distance to $\mathcal{H}_\ell$ to its *head*, the vertex at the shortest distance. Thus we consider P_ℓ as a directed graph. $\mathcal{H}_\ell$ also induces an ordering $(e_{\ell 1}, \ldots, e_{\ell r}, \ldots, e_{\ell p})$, $p = (n-1)n!/2$, of $E(P_\ell)$: edge $e' = v'w'$ comes before edge $e'' = v''w''$ if $\mathrm{dist}(v', \mathcal{H}_\ell) < \mathrm{dist}(v'', \mathcal{H}_\ell)$ or $\mathrm{dist}(v', \mathcal{H}_\ell) = \mathrm{dist}(v'', \mathcal{H}_\ell)$ and $\mathrm{dist}(w', \mathcal{H}_\ell) < \mathrm{dist}(w'', \mathcal{H}_\ell)$. Recall that, in a direct graph, a *source* is a vertex with all the incident edges directed away from it. A *sink* is a vertex with all the incident edges directed towards it.

Proposition 21 (Single source and single sink on R^J_π). *Let* $\mathcal{H}_\ell$, $1 \leq \ell \leq |V(F)|$, *be distinctive for hyperplanes* P_ℓ. *Let also* $J \subseteq \{0, 1, \ldots, n-1\}$ *and* R^J_π *be a* J*-residue of* P^{n-1}_ε. *Then there is a single source and a single sink in* R^J_π.

Proof. It follows from the convexity of the polytope C_π^J corresponding to R_π^J, of Proposition 20. □

The next proposition is the key to the applicability of the permutohedron to our theory, as seen in Proposition 23.

Proposition 22 (Key property). *Let the distances of the vertices of P_ℓ to a hyperplane $\mathcal{H}_\ell$ all be distinct. Denote by $(e_{\ell 1} \dots e_{\ell r} \dots e_{\ell s} \dots e_{\ell p})$ the $\mathcal{H}_\ell$-sequence of $E(P_\ell)$ and by $(v_{\ell 1} \dots v_{\ell r} \dots v_{\ell s} \dots v_{\ell p})$ the sequence of their heads, according to the $\mathcal{H}_\ell$-orientation. Define*

$$\mathcal{J}_{\ell s} = \{e_{\ell r} \mid r \leq s,\ \text{head}(e_{\ell r}) = v_{\ell s}\}, \quad J_{\ell s} = \{\text{color}(e_{\ell r}) \mid e_{\ell r} \in \mathcal{J}_{\ell s}\}.$$

Denote by $R_{\ell s}$ the residue $(v_{\ell s})_{J_{\ell s}}^{P_\ell}$. Then, for $1 \leq s \leq p$, the edge $e_{\ell s}$ is the last edge of the residue $R_{\ell s}$, that is, $\{e_{\ell 1} \dots e_{\ell r} \dots e_{\ell s}\} \cap E(R_{\ell s}) = E(R_{\ell s})$.

Proof. Note that $v_{\ell s}$ is a sink in $R_{\ell s}$, because an edge whose tail is $v_{\ell s}$ must come after $e_{\ell s}$, by definition of $\mathcal{H}_\ell$-sequence. Assume that there exists an edge $e_{\ell t}$ in $R_{\ell s}$ with $t > s$. Take a maximal directed path starting with $e_{\ell t}$ in $R_{\ell s}$. The last vertex v of this path is a sink and we claim that $v \neq v_{\ell s}$. If the tails of $e_{\ell s}$ and $e_{\ell t}$ coincide, then head$(e_{\ell s})$ is closer to $\mathcal{H}_\ell$ than head$(e_{\ell t})$. Therefore $v \neq v_{\ell s}$, otherwise we have that tail$(e_{\ell r})$ is closer to P_ℓ than tail$(e_{\ell t})$, for $e_{\ell r} \in \mathcal{J}_{\ell s}$: if not, then $e_{\ell t}$ would come before $e_{\ell r}$. This implies that $v \neq v_{\ell s}$, since the last edge of the path is not an edge of $\mathcal{J}_{\ell s}$. Thus $R_{\ell s}$ has more than one sink, contradicting Proposition 21. Thus no $e_{\ell t}$ with $t > s$ can exist in $R_{\ell s}$. □

8. The Proof that $\overline{G}\ (\uparrow \Box \downarrow)^*_G\ \overline{H}$

Recall that $F^\star$ is an $(n-1)$-gem inducing S^{n-1}. Let $\overline{F'}$ be the ball complex associated to $F^\star$ in S^{n-1}. With this definition $F^\star$ becomes the 1-skeleton of $\overline{F'}$. Remove a point $\infty \in \overline{F'} \setminus F^\star$ and let $\overline{F} = \overline{F'} \setminus \{\infty\}$. Let $\varepsilon \ll 1$ and P_0 a ε-scaled copy of P_ε^{n-1}, that is, $P_0 = \varepsilon \cdot P_\varepsilon^{n-1}$. Fix a PL homeomorphism of $\overline{F}$ into $\mathbb{R}^{n-1}$, so that the image of the $(N\setminus\{n-1, n\})$-residues are isometric to P_0 and the length of the image of the $(n-1)$-edges are at least 1. We do not distinguish between $F^\star$ and the image of this homeomorphism.

In order to define adequate hyperplanes $\mathcal{H}_\ell$ in $\mathbb{R}^{n-1}$ we fix a *breadth-first search order* (bfs-order, see [3]) to visit the vertices of (the connected graph) F. To find this order use a queue Q for the vertices defined by the following rules: (i) start with Q empty; (ii) choose any vertex v; (iii) put v on Q; (iv) put all the not yet queued v-neighbors in Q (in an arbitrary order) and remove v from Q; (v) if Q is not empty, make $v \leftarrow f$, where f is the first vertex in Q, and go to step (iii); (vi) terminate if Q is empty. The *bsf-order* is the order in which the vertices are enqueued. Fix the notation

$$bfs(F) = (v_1, v_2, \dots, v_\ell, \dots, v_{|V(F)|}).$$

Re-index the permutohedra $P_1, \ldots, P_\ell, \ldots, P_{|V(F)|}$, so that $v_\ell \leftrightarrow P_\ell$, $1 \leq \ell \leq |V(F)|$. From its combinatorial construction, each permutohedron P_ℓ has n $\{0, 1, \ldots, n-2\}$-residues, whose set is denoted by $\mathcal{M}_\ell = \{M_\ell^0, \ldots, M_\ell^{n-1}\}$. For $\ell \in \{2, \ldots, |V(F)|-1\}$, define an injective function $\lambda_\ell : \mathcal{M}_\ell \to \{1, 2, \ldots, |V(F)|\} \backslash \{\ell\}$, as follows. Each vertex of M_ℓ^j is linked by an $(n-1)$-edge to a vertex with the same permutation in a certain $M_k^j \subset P_k \leftrightarrow v_k$, $k \neq \ell$. Define $\lambda_\ell(M_\ell^j) = k$. The function λ_ℓ induces a partition on $\mathcal{M}_\ell$ into two non-empty parts: $\mathcal{M}_\ell^< = \{M_\ell^j \in \mathcal{M}_\ell \mid \lambda_\ell(M_\ell^j) < \ell\}, \mathcal{M}_\ell^> = \{M_\ell^j \mid \lambda_\ell(M_\ell^j) > \ell\}$. In turn, this bipartition permits the definition of an adequate hyperplane $\mathcal{H}'_\ell$ associated to P_ℓ. Let c_ℓ^j be the barycenter of the vertices of M_ℓ^j. Define $d_\ell^<$ as the barycenter of the points c_ℓ^j, for M_ℓ^j in $\mathcal{M}_\ell^<$ and $d_\ell^>$ as the barycenter of the points c_ℓ^j, for M_ℓ^j in $\mathcal{M}_\ell^>$. Define $\mathcal{H}_1$ and $\mathcal{H}_{|V(F)|}$ as arbitrary distinctive hyperplanes for P_1 and $P_{|V(F)|}$ respectively. For $\ell \in \{2, \ldots, V(F)|-1\}$, define $\mathcal{H}'_\ell$ as the hyperplane orthogonal to the line r which links $d_\ell^<$ to $d_\ell^>$ and contains a point d of r, so that $(d, d^<, d^>)$ appear in this order in r and $P_\ell \cap \mathcal{H}'_\ell = \emptyset$. Let $\mathcal{H}_\ell$ be a slight perturbation of $\mathcal{H}'_\ell$ (so that if v is closer to $\mathcal{H}'_\ell$ than w, then v is also closer to $\mathcal{H}_\ell$ than w) so as to become distinctive for P_ℓ. An important aspect of this ordering of the vertices of P_ℓ, induced by the distances to the $\mathcal{H}_\ell$, is that given $1 \leq j_1 < j_2 \leq n-1$ either all the vertices of $M_\ell^{j_1}$ come before the ones of $M_\ell^{j_2}$, or vice-versa. In fact, if we denote by C_{j_1} and C_{j_2} the convex subpolytopes induced by $M_\ell^{j_1}$ and $M_\ell^{j_2}$, it is true that the convex set C_{j_1} is globally either closer to or farther from $\mathcal{H}_\ell$ than C_{j_2}. The crucial point is that the vertices in the elements of $\mathcal{M}_\ell^<$ come before the ones in the elements of $\mathcal{M}_\ell^>$. We say that such a set of hyperplanes $\mathcal{H}_\ell$, $1 \leq \ell \leq |V(F)|$, is *consistent with* $bfs(F)$. See Fig. 17.

At present we have a $bfs(S)$-consistent set of hyperplanes $\mathcal{H}_\ell, 1 \leq \ell \leq |V(F)|$. Fix notation so that the $\mathcal{H}_\ell$-sequence of the edges of P_ℓ is $(e_{\ell 1}, \ldots, e_{\ell r}, e_{\ell, r+1}, \ldots, e_{\ell p})$, where $p = |E(P_0)| = (n-1)n!/2$. Let $(v_{\ell 1}, \ldots, v_{\ell r}, v_{\ell, r+1}, \ldots, v_{\ell p})$ be the corresponding sequence of heads given by the $\mathcal{H}_\ell$-orientation of $E(P_\ell)$, and denote the corresponding sequence of colors by $(c_{\ell 1}, \ldots, c_{\ell r}, c_{\ell, r+1}, \ldots, c_{\ell p})$. Starting with $\overline{G_{00}^\circ} = \overline{G^\circ}$ we produce a sequence of $|V(F)| \cdot p$ pseudogems $\overline{G_{\ell-1,r}^\circ}$, by applying the following operation $o_{\ell-1,r}$ to $\overline{G_{\ell-1,r-1}^\circ}$ and letting $\overline{G_{\ell-1,p}^\circ} = \overline{G_{\ell,0}^\circ}$, for $1 \leq r \leq p$ and $1 \leq \ell \leq |V(F)|$. The operation $o_{\ell-1,r}$, which transforms $\overline{G_{\ell-1,r-1}^\circ}$ into $\overline{G_{\ell-1,r}^\circ}$, is either the $(v'_{\ell r}, c_{\ell r})$-correcting flip, in the case that $e_{\ell r}$ needs to be corrected in $\overline{G_{\ell-1,r-1}^\circ}$, or else the *do nothing operation*, in which case $\overline{G_{\ell-1,r}^\circ} = \overline{G_{\ell-1,r-1}^\circ}$. Note that $\overline{G_{|V(F)|,0}^\circ} = \overline{H}$.

Define $V_\ell^< = \{v \in V(P_\ell) \mid \exists j \mid v \in M_\ell^j \in \mathcal{M}_\ell^<\}$, and let $P_\ell[V_\ell^<]$ denote the subgraph of P_ℓ induced by $V_\ell^<$. The *coboundary* of a subset of vertices W in a graph G is the subset of edges of G having one end in W and the other end in $V(G) \backslash W$. It is denoted by $\delta(W, G)$.

Proposition 23 (True s-flips). *Let $t_\ell = |E(P_\ell)| - |E(P_\ell[V_\ell^<])|$. Then all the operations $o_{\ell-1,r}$, $t_\ell + 1 \leq r \leq p$, are correcting flips. More precisely, they are s-flips.*

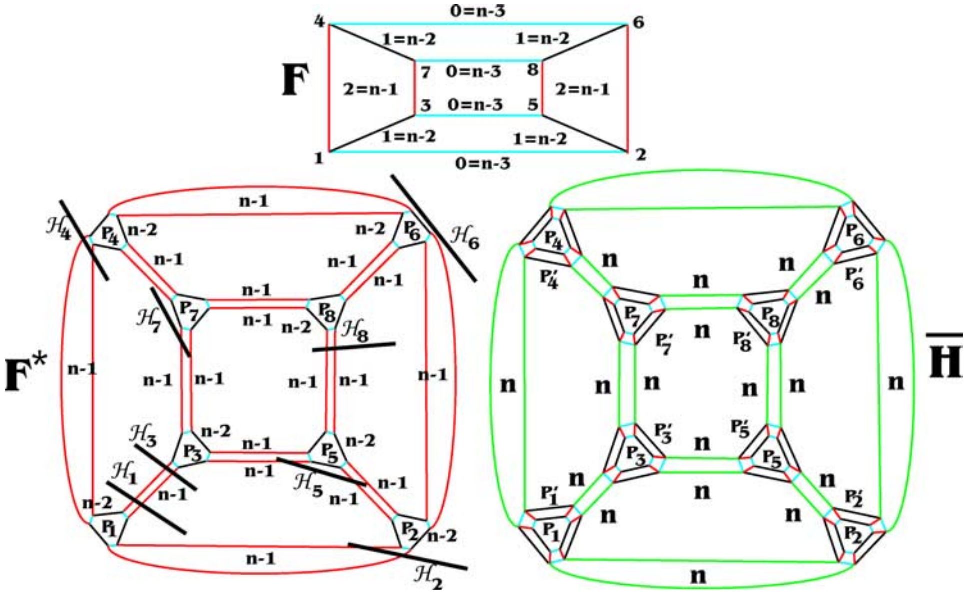

Fig. 17. Example of $bfs(F)$-consistent hyperplanes.

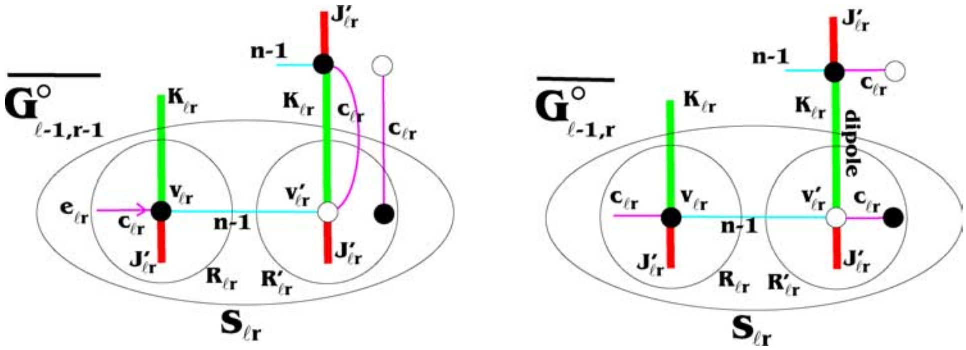

Fig. 18. The true ℓrth dipole slimming.

Proof. Consider first the case $\ell = 1$, $t_\ell = 0$. In this case, no edge of P_ℓ is correct and every $o_{\ell-1,r}$, $1 \le r \le p$, is a correcting flip. Let $J_{\ell r}$ and $R_{\ell r}$ be as defined in Proposition 22. Refer to Fig. 18, where $J'_{\ell r} = J_{\ell r}\backslash\{c_{\ell r}\}$. Define $K_{\ell r} = N\backslash(J_{\ell r} \cup \{n-1\})$. Let $V(R'_{\ell r})$ be the set of primed vertices corresponding to $V(R_{\ell r})$ and $V(S_{\ell r}) = V(R_{\ell r}) \cup V(R'_{\ell r})$. By Proposition 22, $e_{\ell r}$ is the last edge of the $J_{\ell r}$-residue $R_{\ell r}$ of $\overline{G^\circ_{\ell-1,r-1}}$. It follows that, except for the two $c_{\ell r}$-edges involved in the $(v'_{\ell r}, c_{\ell r})$-correcting flip, all the edges in the coboundary of $V(S_{\ell r})$ in $\overline{G^\circ_{\ell-1,r-1}}$ are $K_{\ell r}$-edges. Therefore we may conclude that $(v'_{\ell r})^{\overline{G^\circ_{\ell-1,r}}}_{K_{\ell r}}$ is a dipole, establishing the proposition for $\ell = 1$.

We claim that, for $1 \leq \ell \leq |V(F)|$ and for $h \in \{0, 1, \ldots, n-3\}$, every h-edge of $P_\ell[V_\ell^<]$ is corrected in $\overline{G^\circ_{\ell-1,0}}$. Let $e \in E(P_\ell[V_\ell^<])$ be an h-edge. There exists an $\ell' < \ell$ such that $(v_\ell, \pi), (v_\ell, \pi'), (v_{\ell'}, \pi)$ and $(v_{\ell'}, \pi')$ are all the four vertices of a $\{h, n-1\}$-gon of $F^\star \subset \overline{G}$. By trisecting the two $(n-1)$-edges of this square, we get a corresponding $\{h, n-1\}$-octagon in $\overline{G^\circ}$. Therefore, the s-flip which corrects the h-edge of $P_{\ell'}$ has the side effect of correcting e, breaking the octagon into two squares and proving the claim. This side effect is undesirable, because it makes it impossible, when $\ell > 1$, to perform the correcting flips of $E(P_\ell)$ using the whole sequence $(e_{\ell 1}, \ldots, e_{\ell t_\ell}, e_{\ell, t_\ell+1}, \ldots, e_{\ell p})$ prescribed by the hyperplane $\mathcal{H}_\ell$. However, by Proposition 24 proved below, we can correct the $(n-2)$-edges of $\overline{G^\circ_{\ell-1,0}}$ that are not corrected by a sequence of G-special gem moves. After these corrections have been made, then we get the gem $\overline{G^\circ_{\ell-1,t_\ell}}$, namely, the same n-gem that we would have obtained using the first t_ℓ corrections prescribed by $\mathcal{H}_\ell$. To go from $\overline{G^\circ_{\ell-1,t_\ell}}$ to $\overline{G^\circ_{\ell-1,p}} = \overline{G^\circ_{\ell,0}}$ apply the last $p - t_\ell$ corrections to the subsequence $(e_{\ell,t_\ell+1}, \ldots, e_{\ell p})$ prescribed by $\mathcal{H}_\ell$. The same arguments used for the proof of the case $\ell = 1$ work for this subsequence. Therefore, to finish the proof, it is enough to establish Proposition 24. □

The geometry of the passage $\overline{G^\circ_0} \to \overline{G^\circ_1}$, which *thickens* P_1, is exemplified in Fig. 19 for the case $n = 3$. We show $\overline{G^\circ_{0,2}}$ ($\overline{G^\circ}$ after two s-flips), $\overline{G^\circ_{0,4}}$ ($\overline{G^\circ}$ after four s-flips) and $\overline{G^\circ_{0,6}} = \overline{G^\circ_{1,0}}$ ($\overline{G^\circ}$ after six s-flips). In the example, the 2-permutohedra are hexagons, but our arguments apply to the general dimension n. It helps to think of each short edge in the permutohedra as not a single edge, but a higher dimensional permutohedra with colors $M = \{0, 1, \ldots, n-3\}$, for arbitrary n. The passage $\overline{G^\circ_{\ell-1,0}} \to \overline{G^\circ_{\ell,0}}$ is named the *thickening* of the permutohedron P_ℓ. In this passage, we correct the edges of P_ℓ, and a parallel copy of P_ℓ, denoted P'_ℓ, is formed. If $\ell = 1$, then an extra copy, P''_ℓ, is also produced when we get to $\overline{G^\circ_{1,0}}$. Since P''_ℓ also induces an $(n-2)$-sphere, we modify the homeomorphism, so that it induces an embedded sphere S_1^{n-2}, which separates in $\mathbb{R}^{n-1}$ the thickened P_1 from all the other P_ℓ's.

For $\ell > 1$, in the last flip of the thickening of P_ℓ, $S_{\ell-1}^{n-2}$ breaks into two $(n-2)$-spheres P'_ℓ and S_ℓ^{n-2}. The latter sphere separates the thickened $P_1, \ldots, P_\ell$ from $P_{\ell+1}, \ldots, P_{|V(F)|}$. We might think of the spheres $S_1^{n-2}, S_2^{n-2}, \ldots, S_{|V(F)|-1}^{n-2}$, as perturbations of a moving S_1^{n-2} sphere. Observe that, after the blobs over the $(n-1)$-edges are created, forming $\overline{G^\circ}$, a number of $(n-2)$-spheres induced by the $\{0, 1, \ldots, n-2\}$-residues arise, and all of them have an n-*side* and an $(n-1)$-*side*, because they are crossed transversally by an n-edge followed by an $(n-1)$-edge. Since the flips do not involve the colors $n-1$ and n, these sides are preserved in the process of going from $\overline{G}$ to $\overline{H}$, even though these $(n-2)$-spheres can coalesce and break apart in the process.

Proposition 24 (Final step in proving $\overline{G}$ $(\uparrow \Box \downarrow)^*_G$ $\overline{H}$). *Any $(n-2)$-edge in $P_\ell[V_\ell^<]$ which is not corrected in $\overline{G^\circ_{\ell-1,0}}$ can be corrected by a finite sequence of G-special gem moves.*

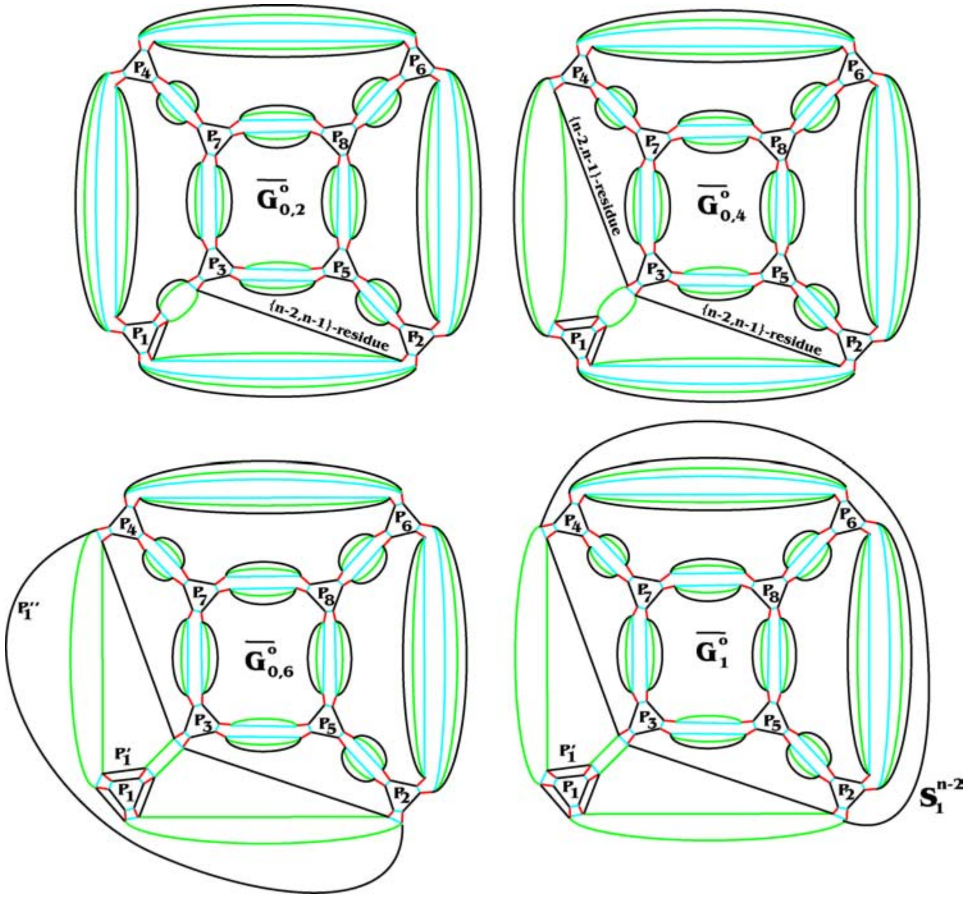

Fig. 19. Thickening the first permutohedron and obtaining a separating sphere S_1^{n-2}.

Proof. Let $M = \{0, 1, \ldots, n-3\}$. We refer to Fig. 22. The $(n-2)$-edges a and b must be flipped to correct an $(n-2)$-edge of permutohedron 15 so that their ends x' and y' become $(n-2)$-adjacent. Let D be the disk attached to the $\{n-2, n-1\}$-residue of x in $\overline{G^\circ}$. The boundary of D meets a cyclic sequence of permutohedra. To illustrate the general case, suppose that these permutohedra are indexed by the bfs-algorithm in F, as shown in Fig. 21. Each permutohedron has a single $(n-1)$-edge in the boundary of G. Consider the sequence of correcting flips from $\overline{G^\circ}$ to $\overline{H}$, which is consistent with the procedure based on the distances to $\mathcal{H}_\ell$ and preserves the order of the permutohedra: namely, if $\ell' < \ell$, then the edges of $P_{\ell'}$ are corrected before the edges of P_ℓ. Then, it follows that the sequence of flips has, as a subsequence, the one displayed in Fig. 21. In our example, the first $(n-2)$-edge of D that is not yet corrected, when correcting the edges of P_ℓ, is (x, y), corrected in the step from Fig. 21(4) to Fig. 21(5). This is because the permutohedron 15

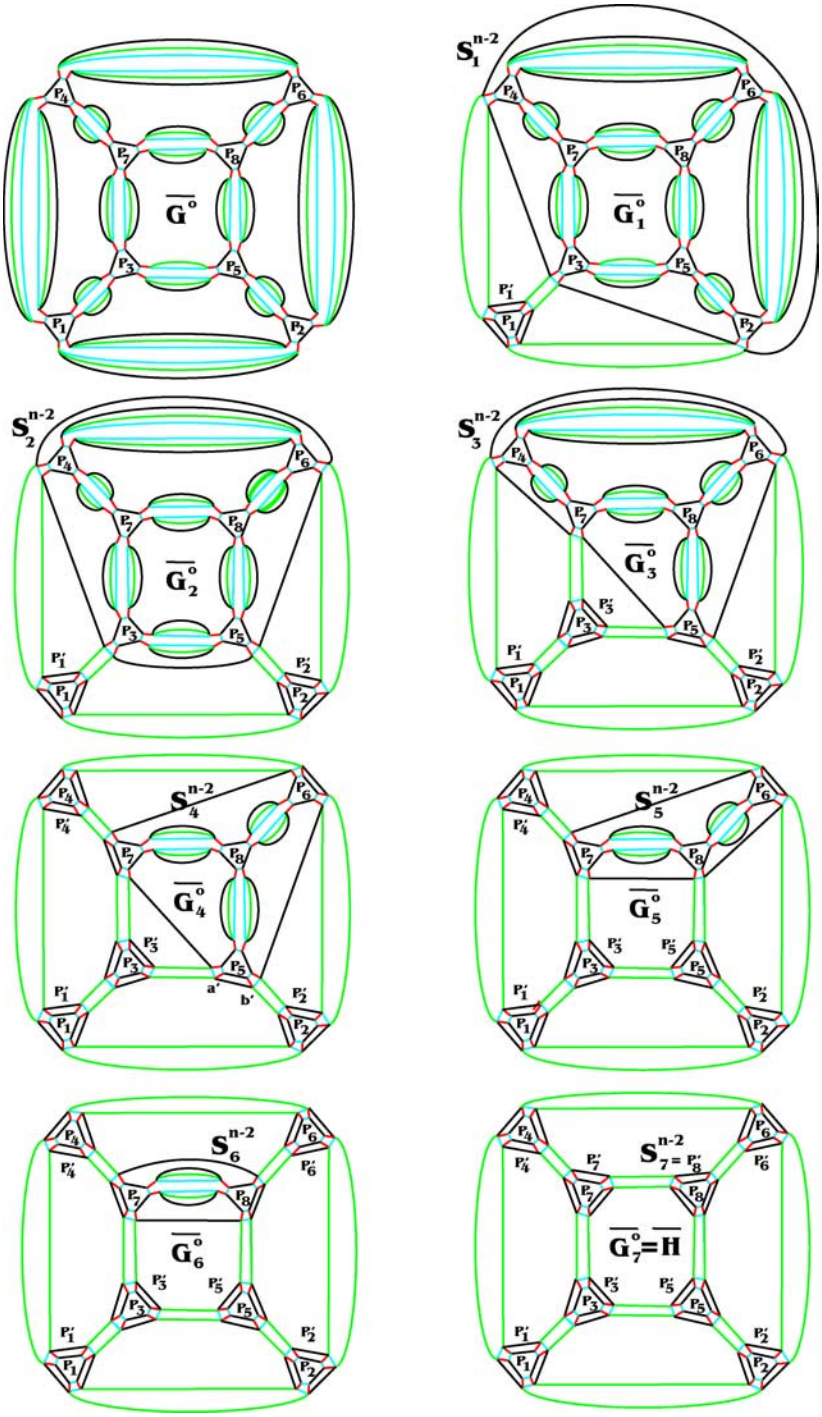

Fig. 20. S_ℓ^{n-2} breaks into P'_ℓ and $S_{\ell+1}^{n-2}$, $1 < \ell \leq |V(F)| - 2$.

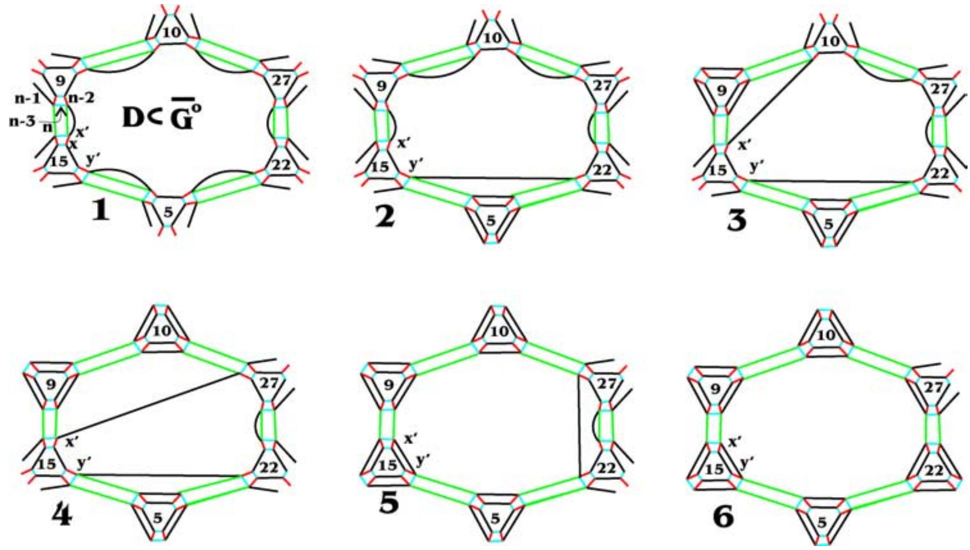

Fig. 21. In an $\{n-3, n-2, n-1, n\}$-residue: a $\{n-2, n-1\}$-disk becomes an $\{n-2, n\}$-disk.

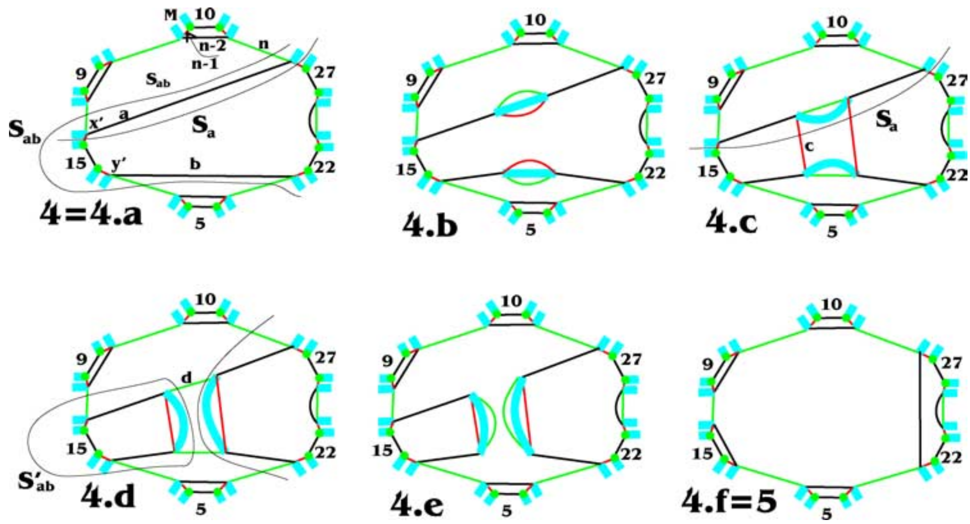

Fig. 22. There exists an $(n-1)$-sphere S_a or else an $(n-1)$-sphere S_{ab} which breaks into two.

is the first one adjacent to two lower indexed permutohedra, namely, 5 and 9. We must then prove that is possible to factorize this correcting flip by G-preserving gem moves.

Refer back to Fig. 22 and consider the $(n-2)$-edges a and b that must be flipped. There are two cases: either a and b belong to distinct $(\{n-2\} \cup M)$-residues or belong to the same $(\{n-2\}\cup M)$-residue. In the former case, we prove that edge c is an $\{n-1\}$-dipole. If the latter case, we prove that edge d is an $\{n\}$-dipole. Suppose the first case. Let S_a be a slight parallel deformation in $\mathbb{R}^{n-1}$ of the $(n-2)$-sphere induced by the $(\{n-2\}\cup M)$-residue of a to its $(n-1)$-side, so that it is crossed only by $(n-1)$-edges. The same sphere S_a is present in part 4.c of Fig. 22, and so we can conclude that c is a $\{n-1\}$-dipole. In the other case, consider S_{ab} be a slight parallel deformation in $\mathbb{R}^{n-1}$ of the $(n-2)$-sphere induced by the $(\{n-2\}\cup M)$-residue of a (and b) to its n-side, so that now it is crossed only by n-edges. Consider part 4.d of Fig. 22, where two $(\{n-1\}\cup M)$-dipoles are created. This modification implies that the $(\{n-2\}\cup M)$-residue breaks into two, and so does the sphere S_{ab}, which breaks into S'_{ab} and S''_{ab}. Since only n-edges cross S'_{ab}, it follows that d is an $\{n\}$-dipole, as we claimed.

Now the conclusion is easy: gems 4.c and 4.d of Fig. 22 are linked by clean flips. Indeed, if c is a dipole, then we can go from gem 4.c to gem 4.d by t-flips that thicken c. If d is a dipole, then we can go from gem 4.d to gem 4.c by t-flips which thicken d. Gem 4.b is obtained from gem 4.a by two blob creations. Gem 4.c is obtained from gem 4.b by a dipole slimming. Gem 4.e is obtained from gem 4.d by a dipole thickening, and gem 4.f is obtained from gem 4.e by two blob cancellations. The proof is complete. □

Summary of the proof of $\overline{G}\ \Box^*_G\ \overline{H}$. By blob creations over the $(n-1)$-edges of $F^\star$ we get $\overline{G^\circ}$. Each one of these blobs consists of an $(n-2)$-sphere, which is embedded in $\mathbb{R}^{n-1}$. The $(n-1)$-edge of the blob links, by a line segment in the interior of this sphere, two opposite points. Each one of the two n-edges incident to the blob extends the $(n-1)$-edge crossing $(n-2)$-sphere. The important point is that after the blobs are created to form $\overline{G^\circ}$, all the original vertices of $\overline{G^\circ}$ are at the $(n-1)$-side of each one of the $(n-2)$-spheres induced by the new $\{0,1,2,\ldots,n-2\}$-residues (each, at creation with two vertices). This property is maintained throughout the correction phase, which in this context means to thicken each one of the permutohedra $P_1, P_2, \ldots, P_\ell, \ldots, P_{|V(F)|}$, in this order. The thickening of permuthedron P_ℓ is accomplished by G-special gem moves, as follows: (1) all the edges of $P_\ell[V_\ell^<]$, except possibly some $(n-2)$-edges, are corrected by s-flips, which correct the edges of previous $P_{\ell'}$, $\ell' < \ell$. Use Proposition 24 to correct, by G-special gem moves, the $(n-2)$-edges of $P_\ell[V_\ell^<]$ which have not yet been corrected. Then correct the edges in $E(P_\ell)\backslash E(P_\ell[V^<])$ by the s-flips, according to the algorithm based on the distances to $\mathcal{H}_\ell$.

9. Conclusion

We have proved the existence of a pair of moves on n-gems, named *gem moves*, which act as a combinatorial counterpart for the homeomorphisms of PL n-manifolds. One of the moves, the *clean flip*, maintains the set of vertices of the gem. The other move, the *blob move*, changes the gem in the simplest possible way in the neighborhood of an edge. To test for homeomorphism between $|G|$ and $|H|$ with $|V(G)| \leq |V(H)|$, the n-manifolds induced by gems G and H, the only difficulty is to find an upper bound for the number $\alpha = \alpha(G, H)$, from which we obtain $\alpha' = \alpha + (|V(H)| - |V(G)|)/2$, and to calculate how many blobs will suffice. The homeomorphism question becomes: is it true that $G^{\alpha'} \;\square^*\; H^{\alpha}$? This question, of course, can be solved in finite time, because there are only finitely many gems equivalent by clean flips, and it is easy to generate them all.

From the theory here developed it follows that, if there is a bound on the number of Pachner moves [16] linking two triangulations of a manifold, then there is a bound for $\alpha(G, H)$. Such a bound on Pachner moves for triangulations of the Seifert fibered manifolds and fibre-free Haken manifolds has recently been produced by Mijatovic [13, 14]. There is a huge gap between the bounds that the theory can currently produce and what one might expect in practice. For instance, the theory of *TS-moves*, the essential part of the computational classification developed in [10], is obtained by allowing the creation of only two blobs, as we show in our final proposition.

Proposition 25 (TS-class implied by two blobs). *Let G be a* 3*-gem and let H be any* 3*-gem in the TS-class* (*cf.* [10]) *of G. Then $G \uparrow^2 G' \;\square^*\; H' \downarrow^2 H$.*

Proof. There are six TS-moves, see [10, pp. 133–137]. In the first 3 we create an $\{i\}$-dipole and cancell another. Clearly for these TS-moves one blob is sufficient: $G \uparrow^1 G' \;\square^*\; H' \downarrow^1 H$. Each one of the other three TS-moves is factored as an $\{i, j\}$-dipole creation, an $\{k\}$-dipole creation, a $\{k\}$-dipole cancelation and an $\{i, j\}$-dipole cancelation. The proposition is a straightforward consequence of this factorization. $\square$

The theory here presented can be made *local* and be used to generalize, for arbitrary dimensions, the results for dimensions 3 and 4 proved in [12]. Another possibly fruitful research project could be the search for new invariants of PL n-manifolds based on the gem moves here introduced. Recently we have shown by means of the computational system BLINK developed by Lauro D. Lins [18] in his thesis that the addition of three blobs followed by flips connect distinct TS-classes in an attractor of a 3-manifold. This avoids the use of U-moves which usually increases by much more than 4 the number of vertices.

References

[1] M. R. Casali, An equivalence criterion for 3-manifolds, *Rev. Mat. Univ. Complutense Madr.* **10** (1997) 129–147.
[2] C. Berge, *Principles of Combinatorics* (Academic Press, 1971).
[3] T. H Cormen, C. E. Leiserson and R. L. Rivest, *Introduction to Algorithms* (MIT Press / McGraw-Hill, 1989).
[4] M. Ferri and C. Gagliardi, Crystallisation moves, *Pacific J. Math.* **100** (1982) 85–103.

[5] M. Ferri, C. Gagliardi and L. Grasselli, A graph-theoretical representation of PL-manifolds. A survey on crystallizations, *Aequationes Math.* **31** (1986) 121–141.

[6] C. Gagliardi, On a class of 3-dimensional polyhedra, *Ann. Univ. Ferrara* **33** (1987) 51–88.

[7] L. Kauffman and S. Lins, *Temperley-Lieb Recoupling Theory and Invariants of 3-Manifolds*, Annals of Mathematical Studies, Vol. 134 (Princeton University Press, Princeton, NJ, 1994).

[8] W. B. R. Lickorish, Simplicial moves on complexes and manifolds, in *Geometry and Topology Monographs, Proceedings of the Kirbyfest*, Vol. 2 (1999), pp. 299–320.

[9] S. Lins and A. Mandel, Graph-encoded 3-manifolds, *Discrete Math.* **57** (1985) 261–284.

[10] S. Lins, *Gems, Computers and Attractors for 3-Manifolds*, Series on Knots and Everything, Vol. 5 (World Scientific, 1995).

[11] S. Lins, Twistors: Bridges among 3-manifolds, *Discrete Math.* **177** (1997) 145–165.

[12] S. Lins and J. S. Carter, Thin-G theory and local moves for gems, *Adv. Math.* **143** (1999) 251–283.

[13] A. Mijatovic, Triangulations of Seifert fibered manifolds, *Math. Ann.* **330** (2004) 235–273.

[14] A. Mijatovic, Triangulations of fibre-free Haken 3-manifolds, *Pacific J. Math.* **219** (2005) 139–186.

[15] C. Rourke and B. Sanderson, *Introduction to Piecewise-Linear Topology* (Springer-Verlag, Berlin, Heildeberg, New York, 1972).

[16] U. Pachner, PL-homeomorphic manifolds are equivalent by elementary shellings, *European. J. Combin.* **12**(1991) 129–145.

[17] A. Schrijver, *Combinatorial Optimization: Polyhedra and Efficiency* (Springer-Verlag, Berlin, Heildeberg, New York, 2003).

[18] L. Lins, BLINK: A new tool for manipulating 3-manifolds, PhD Thesis, under preparation (2006).