University of Southampton

Faculty of Social and Human Sciences

Mathematical Sciences

On the Platonicity of Polygonal Complexes

Yu-Yen Chien

A thesis submitted for the degree of

Doctor of Philosophy

October 2014

UNIVERSITY OF SOUTHAMPTON

ABSTRACT

FACULTY OF SOCIAL AND HUMAN SCIENCES
MATHEMATICAL SCIENCES

Doctor of Philosophy

ON THE PLATONICITY OF POLYGONAL COMPLEXES

by Yu-Yen Chien

In this thesis we study the symmetries of polygonal cell complexes, which are basically 2-dimensional CW-complexes with polygonal faces. In particular we are interested in platonic complexes, of which the automorphism group acts transitively on incident triples of vertex, edge, and face. The main ingredient of the thesis is to study complexes through links, which are graphs describing the local structure of complexes.

We start with discussing how local structure can affect the whole complex. In particular we show that for any symmetric graph, there exists a platonic complex with such link. With computer aid, we find all rigid flexible graphs up to 30 vertices. Such graphs are eligible as links for a classification theorem of platonic polygonal complexes. We also generalize the classification theorem by relaxing the condition on links.

A significant portion of this thesis focuses on products of complexes. We discover the categorical product of polygonal cell complexes, and develop a unique factorization property under some conditions, as well as a concise description of the automorphism groups of products. We also construct another associative product of complexes. These two products interact nicely with two different graph products of link graphs, and can be used to construct platonic complexes with certain links.

In addition to platonic complexes, we construct some complexes with certain type of pathological symmetries. We also study computational problems without involving symmetry, such as the complexity of polygonal complex isomorphism problem, and the decision problem of constructing a complex with prescribed links.

Contents

List of Figures

Author's Declaration

I, Yu-Yen Chien, declare that the thesis entitled *On the Platonicity of Polygonal Complexes* and the work presented in the thesis are both my own, and have been generated by me as the result of my own original research. I confirm that:

- this work was done wholly or mainly while in candidature for a research degree at this University;
- where any part of this thesis has previously been submitted for a degree or any other qualification at this University or any other institution, this has been clearly stated;
- where I have consulted the published work of others, this is always clearly attributed;
- where I have quoted from the work of others, the source is always given. With the exception of such quotations, this thesis is entirely my own work;
- I have acknowledged all main sources of help;
- where the thesis is based on work done by myself jointly with others, I have made clear exactly what was done by others and what I have contributed myself;
- none of this work has been published before submission.

Signed ....

Date ....

Acknowledgements

I would like to express my deep gratitude to my supervisor Professor Ian Leary for all the help and support, mathematically and non-mathematically. Thanks for offering me the opportunity to move from combinatorics to topology, and from the United States to the United Kingdom. Given the topological move is not as trivial as the geographical move, thanks for allowing me so much freedom in research, and having broad interests and insights in my fairly combinatorial work. I would also like to thank all my postgraduate colleagues, who enrich my days in Southampton with stimulating conversations and witticisms over tea cups, coffee mugs, and beer jars. Last but not least, I would like to thank both examiners, Dr. Armando Martino and Dr. Colva Roney-Dougal, for their precious time and valuable advices.

To my parents for the remote but reaching love.

To Pei-Ling for the close and lasting care.

Chapter 1

Introduction

People have always been interested in things with many symmetries. In the three dimensional Euclidean world, the highly symmetric platonic solids are arguably the best examples. These five platonic solids were known by ancient Greeks more than two millenniums ago. Mystically, Timaeus of Locri, an early Pythagorean, associated five platonic solids with four natural elements and the whole universe. Mathematically, platonic solids were treated by Theaetetus of Athens, and in Books XIII to XV of Euclid's Elements [8]. In this thesis, we investigate the platonicity of a generalized version of polyhedra, namely polygonal cell complexes.

Definition 1.1. A **polygonal cell complex** is a 2-dimensional CW-complex satisfying the following extra conditions:

(1) Each 1-cell is an interval of length 1, and each 2-cell is a disc of positive integral circumference.

(2) For a 2-cell of circumference n, the attaching map sends exactly n points evenly distributed on the boundary to the 0-skeleton.

(3) For each boundary segment between the points described in (2), the attaching map sends the segment isometrically onto an open 1-cell.

Intuitively speaking, we can think of each 2-cell as a regular polygon, and the attaching map glues vertices to vertices, and edges to edges. Those 2-cells act like faces of a polyhedron, and we will use the word face to denote a 2-cell alternatively. When a complex is locally finite, by taking the infimum of the lengths of paths joining two points, this induces a complete length metric of the whole complex. A detailed treatment of this metric can be found in Chapter I.7 of [5].

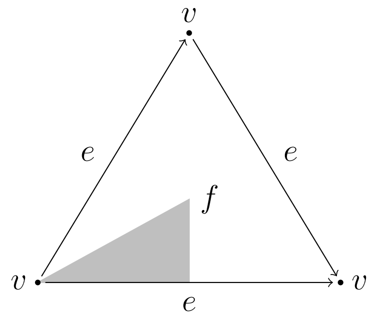

Figure 1.1: a flag of a dunce hat

Note that the attaching map of a face gives a closed walk on the 1-skeleton, with the possibility to visit a vertex more than once. For example, we can have a polygonal cell complex with one vertex, one edge, and one hexagon, where the edge is actually a loop, and the hexagon wraps around the loop six times. Therefore a polygonal cell complex can be quite different from usual polyhedra. To simulate usual polyhedra better, we should place more conditions to rule out pathological cases.

Definition 1.2. A **polygonal complex** is a polygonal cell complex satisfying the following extra conditions:

(1) The attaching map of each cell is injective.

(2) The intersection of any two closed cell is either empty or exactly one closed cell.

Proposition 1.3. A polygonal cell complex X is polygonal if and only if it satisfies the following three conditions:

(1) The 1-skeleton of X is a graph without loops and parallel edges.

(2) The attaching map of each face gives a simple closed cycle on the 1-skeleton.

(3) The intersection of any two faces is either empty, a vertex, or an edge.

Proof. Note that the condition (1) in Definition 1.2 is equivalent to the looplessness in the 1-skeleton plus the condition (2) above. The condition (2) in Definition 1.2 is equivalent to the absence of parallel edges in the 1-skeleton plus the condition (3) above. □

In this thesis, unless otherwise specified, when we use the word complex, it means polygonal cell complex, which may or may not be polygonal. Now we try to define platonicity of complexes. First we define what a flag is for polygonal complexes.

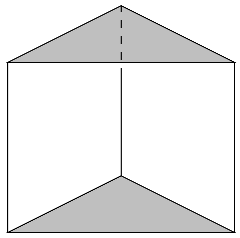

Figure 1.2: transitive on flags, but not platonic

Definition 1.4. For a polygonal complex, a **flag** is a triple (f, e, v) of face, edge, and vertex where f contains e and e contains v. We denote an (f, e, v) flag by fev. A **partial flag** is a face, an edge, a vertex, or an incident pair of the form fe, fv, or ev.

For the case of polygonal cell complexes, the definition of a flag needs to be modified. Take the dunce hat in Figure 1.1 as an example. It has only one vertex, one edge, and one face, but we would like it to have six flags just as a usual triangle. In a polygon, each flag corresponds to a triangle in its barycentric subdivision. We can use this as an alternative definition of a flag, and this definition works for polygonal cell complexes as well. Similarly, a partial flag can be defined as the corresponding 0-cell or 1-cell in the barycentric subdivision. As Figure 1.1 shows, the shaded area is a flag of the dunce hat, and a dunce hat has six flags. This figure to some extent explains why mathematicians choose the word "flag".

For polyhedra, being platonic simply means that the automorphism group acts transitively on flags. However, such a definition does not work perfectly for all polygonal cell complexes. For example, the complex in Figure 1.2 has only the two shaded triangles as its faces. Although the automorphism group acts transitively on fev flags, this complex is not as symmetric as platonic solids. The main drawback is the absence of the transitivity on edges. To exclude such situations and to obtain maximal symmetry, we take the following definition.

Definition 1.5. A polygonal cell complex is said to be **platonic** if its automorphism group acts transitively on flags and on each type of partial flag.

At first glance, it seems very tedious to verify the platonicity of a complex, but in practice we hardly have to worry about this. For a non-degenerate complex, i.e. with at least one face, the platonicity implies that each vertex is incident to an edge, and each edge is incident to a face. If we restrict our discussion to complexes with these two incidence

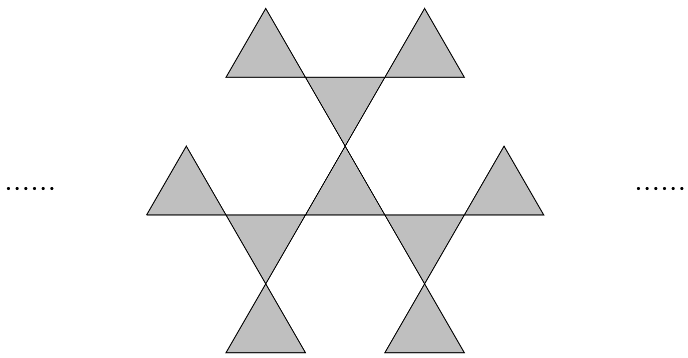

Figure 1.3: a simply-connected platonic complex with disconnected links

conditions, then any partial flag is contained in a flag, and therefore the transitivity on flags implies platonicity. In other words, being platonic and being flag-transitive are equivalent for such complexes.

The main ingredient of this thesis is to study complexes through their local structure. For this purpose, we need a convenient method to describe the neighbourhood of a vertex.

Definition 1.6. For a polygonal cell complex X, the **link** of X at a vertex v is a graph $L(X, v)$ with vertices indexed by ends of edges attached to v, and edges indexed by corners of faces attached to v. Two vertices v_1 and v_2 in $L(X, v)$ are joined by an edge e if and only if the corresponding ends of v_1 and v_2 are joined by the corresponding corner of e.

Basically a link describes the incidence relation of edges and faces at a vertex. Note that $L(X, v)$ can also be identified as the set $\{x \in X \mid d(x, v) = \delta\}$, where d is the distance function in X and δ is some positive number less than 1/2. (In case X is not locally finite, here we can define $d(x, v)$ as $d_c(x, v)$ for any x in a cell c incident to v.) Take Figure 1.1 as an example. Although there is only one edge in the complex, this edge has two ends attached to v, and therefore contributes two vertices to the link at v. Notice that the top corner of the face joins these two ends, and corresponds to an edge joining two vertices in the link at v. The left corner of the face joins the same end of the edge, and hence corresponds to a loop in the link, while the right corner of the face also corresponds to a loop at the other vertex. Therefore the link at v is a graph with two vertices e_1 and e_2, one edge joining e_1 and e_2, and two loops at e_1 and e_2 respectively. For the case of polygonal complexes, it is more straightforward to determine links, and links of polygonal complexes have no loops and parallel edges. For example, the link of the complex in Figure 1.2 is the disjoint union of an edge and a vertex.

In the previous example we encounter a complex with disconnected links. For the discussion of platonic complexes, we would like to exclude complexes with disconnected links. The reason is as follows. Suppose X is a simply-connected platonic complex with disconnected links, as illustrated in Figure 1.3. We can remove all vertices of X, choose a component of the resulting space, and then take the closure of the component to obtain a complex X_0. It is easy to see that X_0 is a platonic complex with connected links, and we can reconstruct X by using X_0 as building blocks. What if X is not simply-connected? Note that the platonicity of a complex can be lifted up to its universal covering. Therefore X can be obtained from the quotient of a simply-connected platonic complex.

To summarize, when we talk about platonic complexes, unless otherwise mentioned, we will assume that each complex is connected, simply-connected, and has connected links. Note that the connectedness of a complex implies that every vertex is incident to an edge, and the connectedness of links implies that every edge is incident to a face. Therefore to verify the platonicity of such complexes, it suffices to verify transitivity on flags. For platonic complexes, we can also assume that there are no loops in the 1-skeleton. With connectedness and transitivity on edges, the existence of a loop implies that the complex has only one vertex, which can be treated separately.

Here we list some terminology and notation which we use through out the thesis. We say a graph is **simple** if there are no loops in the graph, and no parallel edges between any two vertices. The **girth** of a graph is the length of a shortest nontrivial cycle in the graph. The **valency** of a vertex in a graph is the number of ends of edges incident to the vertex, and the **edge valency** of an edge in a complex is the number of sides of faces attaching to the edge. For a graph Γ, we use $V(\Gamma)$ and $E(\Gamma)$ to denote the vertex set and the edge set of Γ respectively, and for a complex X, we use $V(X)$, $E(X)$, and $F(X)$ to denote the vertex set, the edge set, and the face set of X respectively. By X^1 we denote the 1-skeleton of the complex X.

There have been several studies of polygonal cell complexes. Ballman and Brin investigate vertex-transitive complexes under certain curvature conditions, and construct a continuum of non-isomorphic d-gonal complexes with complete link graphs of n vertices for any $d \geq 6$ and $n \geq 3$ in [2]. Świątkowski gives an almost complete description of platonic polygonal complexes with edge valency 3 in [27]. Januszkiewicz, Leary, Valle, and Vogeler classify platonic polygonal complexes with complete graphs as links in [20]. Valle also classifies the case of octahedral link graphs in his PhD thesis [29].

This thesis is composed as follows. In Chapter 2, we investigate various questions related to links. In Chapter 3, we look at rigid flexible link graphs, which have ideal symmetries for the purpose of classifying platonic complexes. In Chapter 4, we try to generalize the classification theorem for CAT(0) platonic polygonal complexes developed in [20]. In Chapter 5, we discuss the tensor product of complexes, which is in fact the categorical product of polygonal cell complexes, and in Chapter 6 we develop another type of product, which interacts nicely with the Cartesian product of link graphs. In Chapters 7 and 8, we discuss the factorization and the symmetry of complexes with respect to the tensor product. In Chapter 9, we study almost platonic complexes, for which the automorphism group acts transitively on each type of partial flag, but not on flags. In Chapter 10, we deal with some computational problems without involving symmetry. We also use GAP to help us find examples and examine properties. All GAP programs are listed in the appendix.

Chapter 2

Link of a Complex

In this chapter we study how the local properties of links can affect the whole complex, and investigate various questions related to links.

If we glue regular d-gons under a regularity condition, namely having two faces meeting at each edge, and m faces meeting at each vertex, then what we can obtain are exactly the five platonic solids. Therefore for polyhedra, platonicity is somehow a byproduct of regularity. Note that the regularity condition above can be rephrased as the link at each vertex being a cycle of length m. When we impose this regularity condition, we have similar results for polygonal cell complexes.

Proposition 2.1. Let X be a finite simply-connected polygonal cell complex. Suppose that each face of X has the same length d, and the link at each vertex is a cycle of length m. Then X is one of the following: tetrahedron, cube, octahedron, dodecahedron, icosahedron, dihedron, hosohedron, which are all platonic.

Proof. Suppose that X has v vertices, e edges, and f faces. The link condition implies that X is actually a surface. Since X is simply-connected, X has Euler characteristic $v - e + f = 2$. By standard counting, we have $v \cdot m = e \cdot 2 = f \cdot d$. We can express v and f in terms of e and rewrite the Euler characteristic equation as

$$\frac{1}{m} + \frac{1}{d} = \frac{1}{2} + \frac{1}{e}.$$

The pair (m, d) has only the following solutions: $(3, 3)$, $(3, 4)$, $(3, 5)$, $(4, 3)$, $(5, 3)$, $(2, e)$, $(e, 2)$. The first five pairs give exactly the five platonic solids, which can be verified by brute force. For $(2, e)$, the complex has e vertices, e edges, and two e-gons as faces. The only possible construction is the so-called dihedron as shown on the left of Figure 2.1.

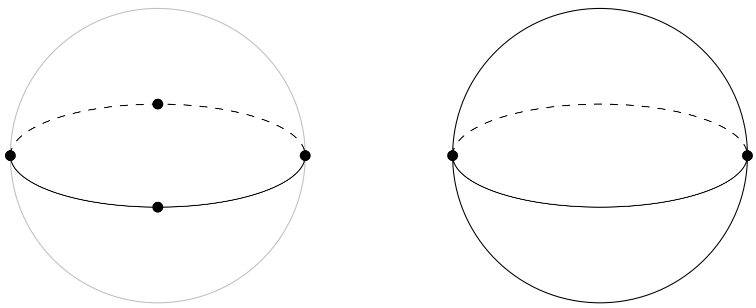

Figure 2.1: (2,4) dihedron and (4,2) hosohedron

For $(e, 2)$, the complex has 2 vertices, e edges, and e 2-gons as faces. The only possible construction is the so-called hosohedron as shown on the right of Figure 2.1. Note that dihedron and hosohedron exist for any positive integer e, even for the extreme case $e = 1$, and dihedron and hosohedron are both platonic. Therefore we finish the proof. □

Now we investigate the finiteness of polygonal cell complexes. Suppose X is a simply-connected polygonal complex such that each face has length at least d, and the girth of the link at each vertex is at least m. If we take one round about a vertex v, then the angle sum is at least $m \cdot \frac{d-2}{d} \cdot \pi$. If $m \cdot \frac{d-2}{d} \cdot \pi \geq 2\pi$, commonly known as the link condition, which can be written as

$$\frac{1}{m} + \frac{1}{d} \leq \frac{1}{2},$$

then X is a CAT(0) space and therefore contractible [5]. The contractibility suggests that X is an infinite complex. In fact for a (m, d) pair satisfying this inequality, one can construct a polygonal complex step after step without having any obstruction, and will always obtain an infinite complex [2]. In what follows, we try to find another criterion forcing infinite complexes. First we need the following lemma.

Lemma 2.2. Suppose X is a finite simply-connected polygonal cell complex. Then the Euler characteristic of X is at least 1.

Proof. Suppose X has v vertices, e edges, and f faces. First we find an arbitrary spanning tree T for the 1 skeleton of X, and then contract T to get a new complex X', which is also simply-connected. Note that T has $v - 1$ edges, and therefore X' has 1 vertex, $e - v + 1$ edges, and f faces. The fundamental group $\pi_1(X')$, a trivial group, can be presented as a group with $e - v + 1$ generators and f relators. Consider the abelianization of $\pi_1(X')$, which is again trivial. Then the presentation can be expressed as f homogeneous

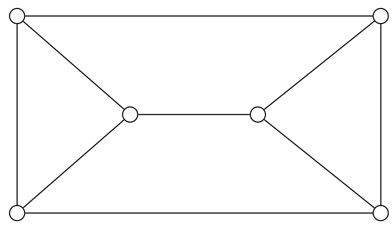

Figure 2.2: a graph of girth 3 and valency 3

equations of $e - v + 1$ unknowns over $\mathbb{Z}$. To have only trivial solution, the number of equations needs to be at least the number of unknowns. So we have $f \geq e - v + 1$, and therefore $v - e + f \geq 1$. $\square$

Proposition 2.3. Let X be a simply-connected polygonal cell complex made of d-gons. Suppose that the link at each vertex has n vertices, and each vertex of the link has valency r. Then X is infinite if we have the following inequality:

$$\frac{1}{n} + \frac{r}{2d} \leq \frac{1}{2}.$$

Proof. We prove this by contradiction. Assume that X is a finite complex satisfying the above inequality, and X has v vertices, e edges, and f faces. By standard counting we have $v = \frac{2e}{n}$ and $f = \frac{r \cdot e}{d}$. Since X is finite and simply-connected, by Lemma 2.2 we have $\frac{2e}{n} - e + \frac{r \cdot e}{d} \geq 1$. Divide both side by $2e$ and then rewrite the inequality to get $\frac{1}{n} + \frac{r}{2d} \geq \frac{1}{2} + \frac{1}{2e} > \frac{1}{2}$, a contradiction. $\square$

Proposition 2.3 and the CAT(0) inequality are implicitly related, as n has a lower bound in terms of girth and valency [4], and the main tool used is the Euler characteristic, which can be viewed as certain curvature condition. Nevertheless, Proposition 2.3 does help to determine some infinite polygonal complexes which are not CAT(0). For example, suppose X is a polygonal complex made of pentagons, and each vertex has the graph in Figure 2.2 as its link. By Proposition 2.3, we know X must be infinite, although X is not CAT(0). The thing is, can we actually build a polygonal complex with such a link? Before answering this question, we first introduce Cayley 2-complexes and Coxeter groups.

Suppose $G = \langle S \mid R \rangle$ is a finitely presented group, where S is the set of generators, and R is the set of relators. Let $\overline{S}$ be the closure of S under inverse, and $\overline{R}$ be the closure of R under inverse and cyclic permutation. The **Cayley 2-complex** $\mathrm{Cay}(G, S, R)$ of this presentation is a polygonal complex with vertex set G. Two vertices $g_1, g_2 \in G$ are joined

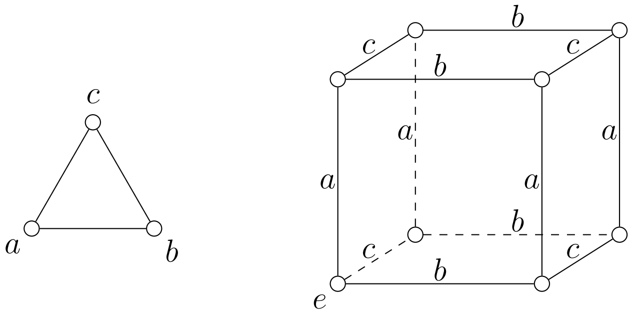

Figure 2.3: K_3 and $\mathrm{Cay}_2(K_3)$

by exactly one edge if and only if there exists $s \in \overline{S}$ such that $g_2 = g_1 * s$. A closed walk $g_1, g_2, \ldots, g_k, g_{k+1} = g_1$ is the boundary of exactly one face if and only if there exists $s_{\alpha_1} s_{\alpha_2} \ldots s_{\alpha_k} \in \overline{R}$ such that $k \geq 3$ and $g_i * s_{\alpha_i} = g_{i+1}$ for all i. Note that we allow no parallel edges and 2-gons in the definition. For example, the complex in Figure 1.3 is the Cayley 2-complex of the group $\langle\, a, b \mid a^3 = b^3 = 1 \,\rangle$. Figure 2.3 gives another example, which is the Cayley 2-complex of $\langle\, a, b, c \mid a^2 = b^2 = c^2 = (ab)^2 = (bc)^2 = (ca)^2 = 1 \,\rangle$.

Note that G acts on $\mathrm{Cay}(G, S, R)$ by left multiplication. Another important property is that $\mathrm{Cay}(G, S, R)$ is always simply-connected. When we have a 1-cycle in $\mathrm{Cay}(G, S, R)$, this represents the trivial element in G, which can be normally generated by relators. This gives us a method to contract the 1-cycle through faces determined by relators.

Let Γ be a finite simple graph with vertex set $\{v_1, v_2, \ldots, v_n\}$. We can define a Coxeter group with respect to Γ for $m \geq 2$ as

$$\mathrm{Cox}_m(\Gamma) = \langle\, v_1, v_2, \ldots, v_n \mid v_1^2, v_2^2, \ldots, v_n^2, (v_i v_j)^m \text{ for adjacent vertices } v_i \text{ and } v_j \,\rangle.$$

Note that we do not treat Γ as the usual Coxeter diagram. A Coxeter group of n generators always has a faithful representation in $GL(n, \mathbb{R})$, where each generator is represented as an involution. A detailed proof can be found in Appendix D of [9]. In particular, the group $\mathrm{Cox}_m(\Gamma)$ is nontrivial, and therefore we can construct the corresponding Cayley 2-complex, denoted by $\mathrm{Cay}_m(\Gamma)$, as illustrated in Figure 2.3.

Proposition 2.4. For any finite simple graph Γ, $\mathrm{Cay}_m(\Gamma)$ is a polygonal complex.

Proof. Let Γ be a finite simple graph with vertex set $\{v_1, v_2, \ldots, v_n\}$. We verify the three conditions in Proposition 1.3 to show the polygonality of $\mathrm{Cay}_m(\Gamma)$. The condition (1) is directly from the definition. The conditions (2) and (3) are related to the word problem for Coxeter groups. According to Tits' solution [28], for a Coxeter group $G = \langle\, v_1, v_2, \ldots, v_n \mid$

$v_1^2, v_2^2, \ldots, v_n^2, (v_iv_j)^{m_{ij}}\rangle$, a word in $\{v_1, v_2, \ldots, v_n\}$ is trivial in G if and only if it can be reduced to the empty word through a sequence of the following two operations:

(i) Delete a subword of the form (v_i, v_i).

(ii) Replace an alternating subword of the form $(v_i, v_j, \ldots)$ of length m_{ij} by the alternating word $(v_j, v_i, \ldots)$ of the same length m_{ij}.

The condition (2) is essentially saying that for any adjacent v_i and v_j in Γ, $(v_iv_j)^m$ has no proper nonempty subword which is trivial in $\mathrm{Cox}_m(\Gamma)$. As for the condition (3), it suffices to show that there are no two faces meeting at two non-adjacent vertices. This is essentially saying that for any adjacent pairs of vertices (v_i, v_j) and $(v_{i'}, v_{j'})$ in Γ, an alternating word in v_i and v_j of length l, where $2 \leq l \leq 2m-2$, followed by an alternating word in $v_{i'}$ and $v_{j'}$ of length l', where $2 \leq l' \leq 2m-2$, is always nontrivial in $\mathrm{Cox}_m(\Gamma)$. It is straightforward to verify these two statements by Tits' solution, and hence we know that $\mathrm{Cay}_m(\Gamma)$ is a polygonal complex. □

Proposition 2.5. For any finite simple graph Γ, there is a simply-connected vertex-transitive polygonal complex such that the link at each vertex is isomorphic to Γ, and each face is a polygon with an even number of sides.

Proof. Suppose that Γ has vertex set $\{v_1, v_2, \ldots, v_n\}$. Consider the Cayley 2-complex $\mathrm{Cay}_m(\Gamma)$. By definition, the trivial element e has $\{v_1, v_2, \ldots, v_n\}$ as its neighbours. Note that v_i, e, v_j form a corner of a face in $\mathrm{Cay}_m(\Gamma)$ if and only if $(v_iv_j)^m = 1$, if and only if v_i and v_j are adjacent in Γ. Therefore the link of $\mathrm{Cay}_m(\Gamma)$ at e is isomorphic to Γ. Moreover, the action of $\mathrm{Cox}_m(\Gamma)$ on $\mathrm{Cay}_m(\Gamma)$ implies that the link at each vertex is isomorphic. Note that $\mathrm{Cay}_m(\Gamma)$ is simply-connected, vertex-transitive, and polygonal by Proposition 2.4. Also notice that each face has $2m$ sides. □

Similarly to complexes, we can define a flag of a graph as an edge in its barycentric subdivision. For a simple graph, a flag can also be viewed of as an incident pair of edge and vertex. A graph is said to be **symmetric** if its automorphism group acts transitively on vertices, edges, and flags. For a platonic complex, it is easy to see that its link is symmetric. Therefore to obtain a platonic complex, we at least need to assume that the link of the complex is symmetric.

Proposition 2.6. For any finite simple symmetric graph Γ, there is a simply-connected platonic polygonal complex such that the link at each vertex is isomorphic to Γ.

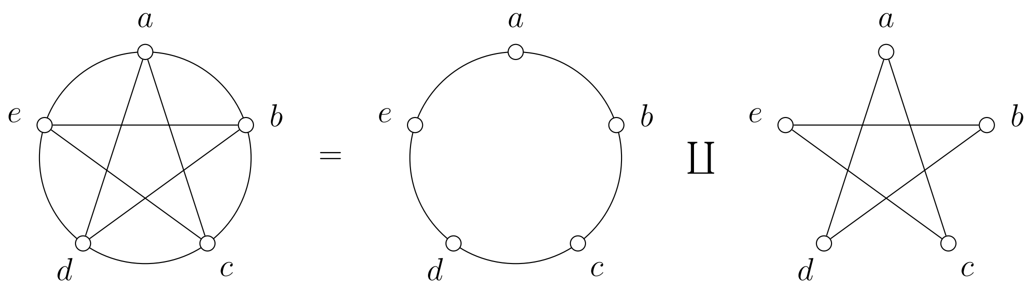

Figure 2.4: a closed decomposition of K_5

Proof. Suppose that Γ has vertex set $\{v_1, v_2, \ldots, v_n\}$. Consider the Cayley 2-complex $\mathrm{Cay}_m(\Gamma)$. We know that $\mathrm{Cay}_m(\Gamma)$ is vertex-transitive, simply-connected, and polygonal, so it suffices to show that it is transitive on flags containing the trivial element e. We use (e, v_i, v_j) to denote the flag of vertex e, the edge between e and v_i, and the face with the corner v_i, e, v_j. For any two flags (e, v_i, v_j) and (e, v_k, v_l), (v_i, v_j) and (v_k, v_l) represent two directed edges in Γ. Since Γ is symmetric, there exists a graph automorphism mapping (v_i, v_j) to (v_k, v_l), and such automorphism determines an automorphism on $\mathrm{Cay}_m(\Gamma)$, mapping (e, v_i, v_j) to (e, v_k, v_l). Therefore $\mathrm{Cay}_m(\Gamma)$ is flag-transitive. □

Remark. Note that in the construction of $\mathrm{Cay}_m(\Gamma)$, we can choose different m's for different relators $(v_i v_j)^m$'s, and we still obtain a vertex-transitive complex. However, such construction only gives complexes with polygons of even length. For any given (platonic) graph Γ, can we always find a vertex-transitive (platonic) complex such that each link is isomorphic to Γ, and each face is an odd polygon? We do not know the answer yet.

We introduce another construction, which has less freedom on the length of polygons, but with the possibility to have some odd polygons. For a graph Γ, a **closed decomposition** $\mathcal{D}$ of Γ is a family of subgraphs of Γ such that each subgraph is a closed walk, and the edge set of Γ is the disjoint union of the edge set of subgraphs in $\mathcal{D}$. Figure 2.4 is an example of closed decomposition $\mathcal{D} = \{abcde, acebd\}$ of the complete graph K_5, where $\{abcedacebd\}$ is another closed decomposition. Note that we express a closed walk as a word of vertices, and in the expression we do not repeat the starting vertex when we finish a closed walk. By arguably the first theorem in graph theory due to Euler [3], a graph admits a closed decomposition if and only if it admits a closed decomposition of cardinality 1, if and only if the valency of each vertex is even. Graphs satisfying these equivalent conditions are commonly known as Eulerian graphs.

Suppose a graph Γ has vertex set $\{v_1, v_2, \ldots, v_n\}$, and $\mathcal{D} = \{w_1, w_2, \ldots, w_m\}$ is a closed decomposition of Γ, where each w_i is a word of $\{v_1, v_2, \ldots, v_n\}$. Define the group

$$G(\Gamma, \mathcal{D}) = \langle\, v_1, v_2, \ldots, v_n \mid v_1^2, v_2^2, \ldots, v_n^2, w_1, w_2, \ldots, w_m \,\rangle$$

and denote the Cayley 2-complex of this presentation by $\mathrm{Cay}(\Gamma, \mathcal{D})$. Note that the group $G(\Gamma, \mathcal{D})$ might be trivial. Under certain small cancellation conditions, the small cancellation theory over free products [22] helps to ensure nontriviality. In the extreme case $m = 1$, which we can manage to find, it is obvious that $G(\Gamma, \mathcal{D})$ is nontrivial.

Proposition 2.7. Suppose Γ is a finite simple graph with closed decomposition $\mathcal{D}$. If $G(\Gamma, \mathcal{D})$ is nontrivial, then the link of $\mathrm{Cay}(\Gamma, \mathcal{D})$ is isomorphic to Γ.

Proof. Suppose Γ has vertex set $\{v_1, v_2, \ldots, v_n\}$, and $\mathcal{D} = \{w_1, w_2, \ldots, w_m\}$. By definition, the trivial element e has $\{v_1, v_2, \ldots, v_n\}$ as its neighbours in $\mathrm{Cay}(\Gamma, \mathcal{D})$. Note that v_i, e, v_j form a corner of a face in $\mathrm{Cay}(\Gamma, \mathcal{D})$ if and only if v_iv_j or v_jv_i is a piece of some w_k, if and only if v_i and v_j are adjacent in Γ. Hence the link of $\mathrm{Cay}(\Gamma, \mathcal{D})$ at e is isomorphic to Γ. By the property of Cayley 2-complexes, the link at each vertex is isomorphic to Γ. $\square$

Remark. The construction $\mathrm{Cay}(\Gamma, \mathcal{D})$ gives us a vertex-transitive complex. However, when we assume that Γ is symmetric, unlike $\mathrm{Cay}_m(\Gamma)$, the automorphism of Γ does not always incude an automorphism of $\mathrm{Cay}(\Gamma, \mathcal{D})$. Unless the decomposition is highly symmetric, $\mathrm{Cay}(\Gamma, \mathcal{D})$ is not platonic in general.

Definition 2.8. A simple graph is **rigid** if an automorphism fixing a vertex and all its neighbours must be trivial. Similarly, a polygonal complex is **rigid** if an automorphism fixing a vertex and all its neighbours must be trivial.

Proposition 2.9. Suppose that X is a connected polygonal complex. If the link at each vertex of X is rigid, then X is rigid.

Proof. Let ϕ be an automorphism of X fixing a vertex v and all its neighbours. We want to show that ϕ fixes each vertex of X by induction on the distance of the vertex to v. By the assumption, each vertex of distance 1 to v is fixed. Suppose that ϕ fixes each vertex within distance n to v, and w is a vertex of distance $n+1$ to v. Take a geodesic from v to w, and look at the last two vertices x and y before arriving at w, as illustrated in Figure 2.5. Note that all neighbours of y are within distance n of v, and therefore are fixed by ϕ.

Suppose f is a face containing the edge $\{x, y\}$. Note that f must have a corner (x, y, z_i), where z_i is a neighbour of y. Since X is a polygonal complex, the action of ϕ on

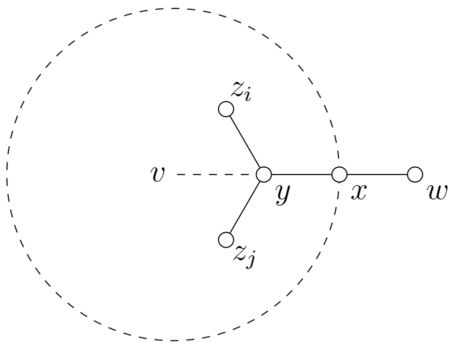

Figure 2.5: a ball of radius n centred at v

f is determined by its action on the corner (x, y, z_i). Therefore ϕ fixes f, and all other faces incident to $\{x, y\}$. In the link of x, the above statement is essentially saying that the induced action of ϕ fixes every edge incident to the vertex $\{x, y\}$. By the rigidity of links, ϕ acts trivially on the link of X. In particular, ϕ fixed the vertex w in the complex X. By induction and the connectedness of X, we know that ϕ fixes each vertex of X. Since X is a polygonal complex, ϕ must be trivial, and hence X is rigid. □

Remark. In the proof, if we only assume that ϕ fixes v, a neighbour v' of v, and all other neighbours of v forming a corner with v and v', then we can still derive the triviality of ϕ by the same argument. Therefore, X has slightly stronger rigidity than our definition.

Corollary 2.10. Suppose that X is a locally finite connected polygonal complex where the link at each vertex is rigid. Then for any vertex v, edge e, and face f, the stabilizer G_v, G_e, G_f are finite subgroups of the automorphism group G of X.

Proof. For an arbitrary vertex v, we can assume that v has n neighbours since X is locally finite. By Proposition 2.9, we know that X is rigid, and an element of G_v is completely determined by its action on the neighbours of v. Therefore the order of G_v is less than or equal to $n!$. Then G_{ev} and G_{fv} are finite as they are subgroups of G_v. Note that G_e and G_f are finite index supergroups over G_{ev} and G_{fv} respectively. Hence G_v, G_e, G_f are all finite subgroups of G. □

Chapter 3

Rigid Flexible Links

In [20], a classification theorem of CAT(0) platonic polygonal complexes is developed for any link graph Γ with the following two properties:

(1) Γ is rigid.

(2) Γ is vertex-transitive, and the stabilizer of a vertex v can permute neighbours of v arbitrarily (namely acts as the symmetric group on neighbours of v).

In this chapter, we investigate and give some examples of graphs with the above two properties. To be concise, we make the following definition.

Definition 3.1. A graph is **flexible** if it satisfies the second condition above.

To start with, complete graphs K_n are the most trivial examples, and cycles obviously have these two properties as well. Note that K_n can be characterized as a simple graph of n vertices where each vertex is of valency $n-1$, and a cycle can be characterized as a connected simple graph where each vertex is of valency 2.

For $n \geq 3$, consider a bipartite graph $B_{n,n}$ such that each partite set has n vertices, and each vertex is of valency $n-1$. For each vertex v in $B_{n,n}$, we can find a unique nonadjacent vertex w in the other partite set, whereas for w the unique nonadjacent vertex in the other partite set is v. This gives us a way to pair vertices in $B_{n,n}$, and then determine edges completely. In other words, $B_{n,n}$ is unique up to isomorphism. Figure 3.1 is a particular drawing of $B_{5,5}$, where partite sets and vertex pairs can be easily recognized. Note that any isomorphism of $B_{n,n}$ preserves such pairing, and the rigidity of $B_{n,n}$ follows straightforwardly. It is not hard to see that the automorphism group of $B_{n,n}$ is $C_2 \times S_n$, which acts transitively on vertices of $B_{n,n}$. Moreover, the stabilizer of a vertex v can permute the $n-1$ neighbours of v arbitrarily. Hence $B_{n,n}$ is flexible.

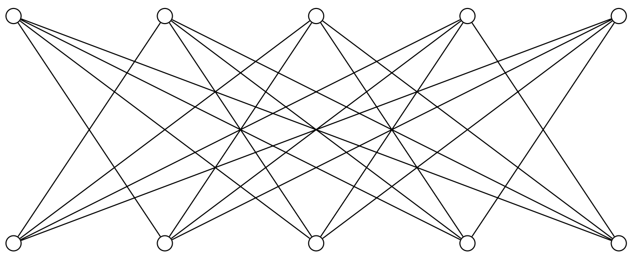

Figure 3.1: a drawing of $B_{5,5}$

Another example is the hypercube graph Q_n, namely the 1-skeleton of the n-dimensional hypercube. We can also define Q_n as a simple graph with vertex set $\{(\delta_1, \dots, \delta_n), \delta_i \in \{0, 1\}\}$, and two vertices are adjacent if and only if they take different values at exactly one coordinate. There are two standard types of automorphisms of Q_n. The first type comes from permuting the n coordinates, and hence can arbitrarily permute the n neighbours of the origin $(0, \dots, 0)$. The second type is generated by swapping the value at a given coordinate, and hence can map the origin to any other vertex. Note that these two types of automorphisms generate $C_2 \wr S_n$, and therefore Q_n is flexible. For the rigidity, note that the vertices of distance i to the origin are exactly the vertices with i nonzero coordinates, and they have distinct sets of neighbours of distance $i - 1$ to the origin. When an automorphism fixes the neighbours of the origin, it fixes vertices of distance 2, and therefore vertices of distance 3, and so on. This shows that Q_n is rigid. By the orbit-stabilizer theorem, Q_n has $2^n \cdot n!$ automorphisms, as many as the order of $C_2 \wr S_n$. Therefore the automorphism group of Q_n is exactly $C_2 \wr S_n$.

Examples of rigid flexible graphs other than these are not immediately obvious. Suppose that Γ is a graph with n vertices satisfying the above two conditions, and G is the automorphism group of Γ. Note that G is a transitive permutation group on n vertices, and the vertex stabilizer G_v is isomorphic to the symmetric group S_m, where m is the number of neighbours of v. Conversely, to construct such a graph, we start with a transitive group G on $\{1, \dots, n\}$ such that the stabilizer G_n is isomorphic to S_m, and there is a suborbit N of length m under the action of G_n. The induced action of S_m on N gives a quotient of S_m, which is either S_m, C_2, or trivial, except when $n = 4$ it could be S_3, the quotient by Klein 4-subgroup. To act transitively on a suborbit of length m, the order of the quotient is a multiple of m, and the only possible quotient is S_m. In other words, G_n acts as a symmetric group on N, which should be viewed as neighbours of n. By choosing these edges and their image under the action of G, we construct a graph

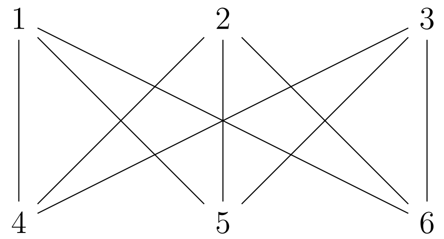

Figure 3.2: a complete bipartite graph $K_{3,3}$

Γ with vertex set $\{1, \ldots, n\}$ and edge set $\{\, g(xn) \mid x \in N,\, g \in G \,\}$. The construction guarantees that G is a subgroup of $\mathrm{Aut}(\Gamma)$, and therefore Γ is flexible. Note that Γ is rigid if and only if G is exactly $\mathrm{Aut}(\Gamma)$. It is possible that the construction creates extra automorphisms other than G, and we need to rule out such cases. For example, consider the group $G = \langle\, (1,2,3), (1,2)(4,5), (1,4)(2,5)(3,6) \,\rangle$, a group of order 36. The stabilizer G_6 is generated by $(1,2,3)$ and $(1,2)(4,5)$, and acts transitively on $\{1,2,3\}$ as S_3. The above construction gives the complete bipartite graph $K_{3,3}$ in Figure 3.2. The automorphism group of $K_{3,3}$ is $C_2 \times S_3 \times S_3$, a group of order 72, and $K_{3,3}$ is not rigid.

We write a GAP program to take care of the above task for $n \leq 30$. Program 1 in the appendix defines a function $\mathrm{Link}(n)$, which tests all transitive groups on $\{1, \ldots, n\}$. Let G be the (n, i)-transitive group, namely the i-th transitive group on $\{1, \ldots, n\}$ in GAP library. We need to verify G_n is isomorphic to S_m for some suborbit of length m, and the resulting Γ has an automorphism group of order $|G|$. In the program, we exclude the cases $m = 2$ and $m = n - 1$, as they give exactly cycles and complete graphs. Since computing the automorphism group of a graph is no harder than verifying permutation group isomorphism [1], initially we only require G_n to have order $m!$, then pick up cases when $\mathrm{Aut}(\Gamma)$ is of order $|G|$, and verify the isomorphism between G_n and S_m in the last stage. We test all transitive groups up to $n = 30$, the limit of GAP library to date. We have a coincidence that in the last stage G_n is always isomorphic to S_m. Can we conclude such an isomorphism simply from G_n having order $m!$ for some suborbit of length m? The $(12, 43)$-transitive group gives a counter example. The stabilizer G_{12} is of order 6, and G_{12} has a suborbit $\{3, 6, 9\}$. However G_{12} is actually isomorphic to C_6, acting unfaithfully on $\{3, 6, 9\}$. In Table 3.1, we list all rigid flexible graphs up to 30 vertices, with cycles, complete graphs, and $B_{n,n}$ omitted for concision. For rows with two suborbits for N, two different suborbits yield isomorphic graphs.

Transitive group	Suborbit N	Group structure	Graph description
$(14, 16)$	$\{3, 5, 7, 13\}$	$\mathrm{PSL}(3, 2) : C_2$	isomorphic to Figure 3.3
$(16, 190)$	$\{7, 9, 12\}$ $\{8, 10, 11\}$	$\mathrm{GL}(2, 3) : C_2$	Figure 3.4
$(16, 748)$	$\{7, 9, 12, 14\}$ $\{8, 10, 11, 13\}$	$C_2 \wr S_4$	hypercube graph Q_4
$(16, 1328)$	$\{3, 4, 6, 10, 11\}$	$((C_2 \times C_2 \times C_2 \times C_2) : A_5) : C_2$	antipodal quotient of Q_5
$(20, 36)$	$\{7, 9, 11\}$ $\{8, 10, 12\}$	$C_2 \times A_5$	1-skeleton of dodecahedron
$(24, 281)$	$\{12, 13, 17\}$	$S_3 \times S_4$	Figure 3.5
$(28, 80)$	$\{5, 10, 14, 26\}$ $\{6, 9, 13, 25\}$	$C_2 \times (\mathrm{PSL}(3, 2) : C_2)$	double cover of $(14, 16)$
$(30, 178)$	$\{2, 7, 20, 25\}$	$S_5 \times S_3$	

Table 3.1: rigid flexible graphs, all bipartite except $(16, 1328)$ and $(20, 36)$

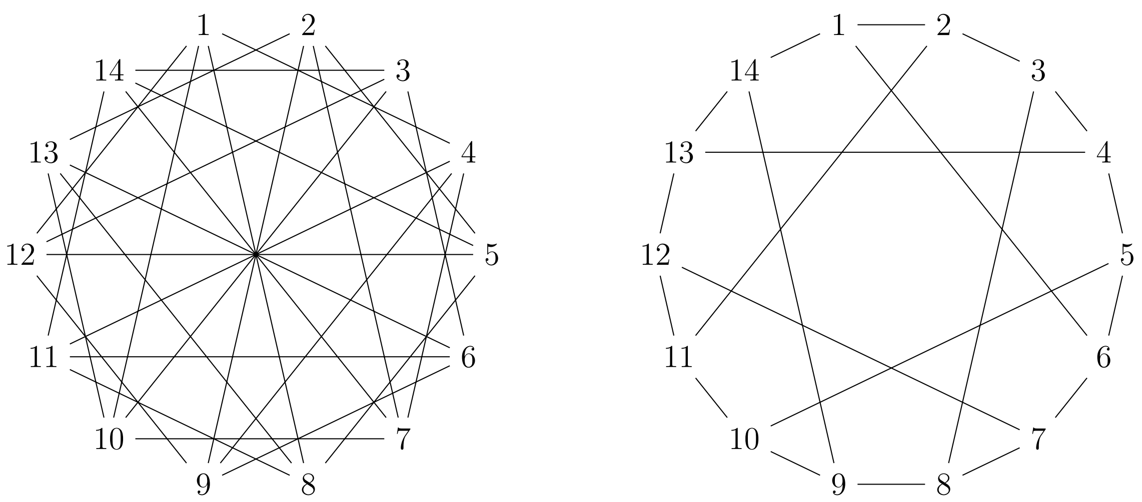

Figure 3.3: a graph isomorphic to the $(14, 16)$-graph and the Heawood Graph

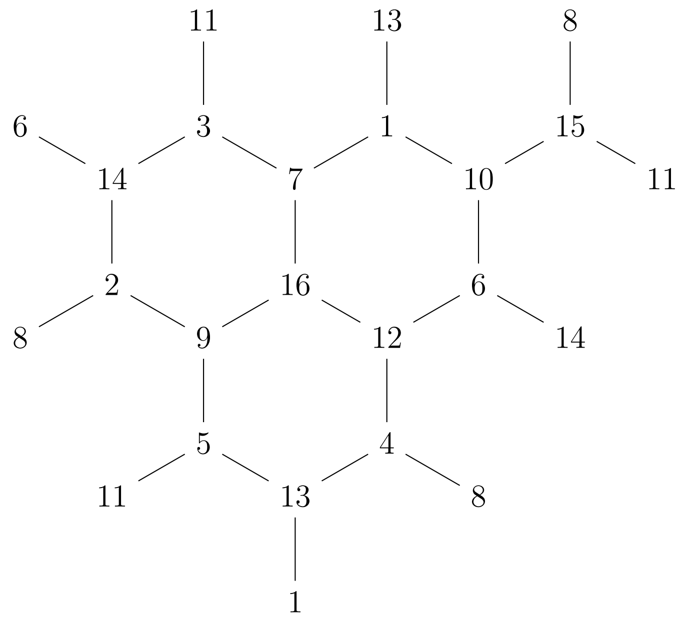

Figure 3.4: the $(16, 190)$-graph

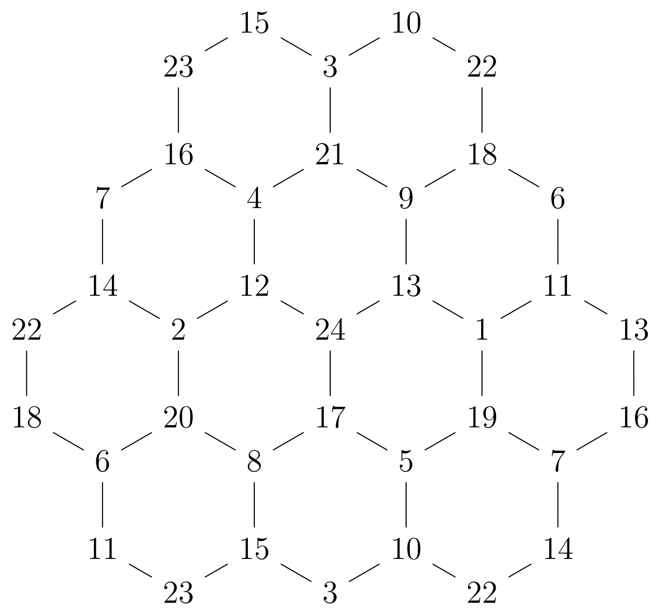

Figure 3.5: the $(24, 281)$-graph

The graph given by the $(14, 16)$-transitive group can be obtained from the incidence relation in the Fano plane. It is a bipartite graph of 14 vertices, representing 7 points and 7 lines of the Fano plane, and we join two vertices if they are non-incident pair of point and line in the Fano plane. Alternatively, if we join two vertices if they are incident pair of point and line, this gives the so-called Heawood graph [16]. In the Heawood graph, we can join vertices at distance 3 by edges and then remove all the original edges. This also gives the $(14, 16)$-graph, as illustrated in Figure 3.3. Program 2 in the appendix verifies that the $(14, 16)$-graph is isomorphic to the graph on the left of Figure 3.3, and the permutation $(2, 4, 8)(3, 11, 7)(5, 9)(6, 10, 14)$ of the $(14, 16)$-graph gives an isomorphism, mapping neighbours $\{3, 5, 7, 13\}$ of 14 to neighbours $\{11, 9, 3, 13\}$ of 6 in Figure 3.3.

Judging from the group structures, it is reasonable to speculate that the $(28, 80)$-graph is related to the $(14, 16)$-graph. In our program, the $(28, 80)$-graph has vertex set $\{1, 2, \dots, 28\}$, and for any odd vertex n, the stabilizer of n also stabilizes $n + 1$. To merge these stabilized pairs, we define a map sending each vertex n to the vertex $\lceil \frac{n}{2} \rceil$, and each edge (n, m) to the edge $(\lceil \frac{n}{2} \rceil, \lceil \frac{m}{2} \rceil)$. The resulting simple graph is isomorphic to the $(14, 16)$-graph, which is verified in Program 2. The program also verifies that the map defined above is actually a **double cover**, namely a surjective locally homeomorphic map such that each vertex (edge) has two vertices (edges) as its preimage.

The graph given by the $(16, 1328)$-transitive group is of particular interest. It can be obtained by identifying antipodal vertices in Q_5. In a hypercube graph, two vertices are antipodal if they take different values at each coordinate. Note that two vertices are adjacent in Q_n if and only if their antipodal vertices are adjacent. Therefore we can well define the antipodal quotient Q_n^* of Q_n. Notice that an automorphism of Q_n preserves antipodal relations, and hence induces an automorphism of Q_n^*. By the flexibility of Q_n, the induced action of $\mathrm{Aut}(Q_n)$ on Q_n^* shows that Q_n^* is flexible.

To discuss rigidity, we introduce another way to describe Q_n^*. For each pair of antipodal vertices in Q_n, we choose the vertex with value 0 in the n-th coordinate to be their representative. These representatives form a copy of Q_{n-1} inside Q_n. For a vertex $v = (\delta_1, \dots, \delta_{n-1}, 0)$, it has $n - 1$ neighbours in Q_{n-1}, and the last neighbour $(\delta_1, \dots, \delta_{n-1}, 1)$ can be identified in Q_n^* with $(1 - \delta_1, \dots, 1 - \delta_{n-1}, 0)$, the antipodal vertex of v in Q_{n-1}. That is to say, if we join each pair of antipodal vertices in Q_{n-1}, the resulting graph $\overline{Q_{n-1}}$ is isomorphic to Q_n^*. For example, $Q_4^* \cong \overline{Q_3}$ is isomorphic to $K_{4,4}$, which is not rigid. For higher dimensional cases, we have the following result.

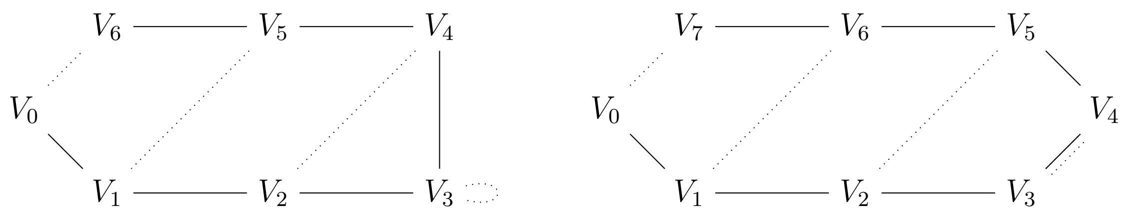

Figure 3.6: structures of $\overline{Q_6}$ and $\overline{Q_7}$

Proposition 3.2. For $n \geq 4$, $\overline{Q_n}$ is rigid and flexible.

Proof. It suffices to show that $\overline{Q_n}$ is rigid. Let $V_i = \{(\delta_1, \ldots, \delta_n) \in \overline{Q_n} \mid \sum_{j=1}^{n} \delta_j = i\}$. Note that for every vertex in V_i, we can find its i neighbours in V_{i-1} and $n-i$ neighbours in V_{i+1} through swapping the value at one coordinate, and find the last neighbour in V_{n-i}, the antipodal one. Figure 3.6 illustrates the structure of $\overline{Q_n}$, where solid edges mean value-swapping and dotted edges represent antipodal pairs. For $1 \leq i \leq \lfloor \frac{n}{2} \rfloor - 1$, each vertex in V_{i+1} has $i+1$ distinct neighbours in V_i, and each vertex in V_{n-i} has a unique neighbour in V_i. In particular, if V_{i-1} and V_i are fixed pointwise, then so are V_{i+1} and V_{n-i}. Consequently, when we fix the origin $(0, \ldots, 0)$ and its neighbours, namely all vertices in V_0, V_1, and V_n, by induction we can fix all other V_j's pointwise, with the only exception being V_{k+1} when $n = 2k+1$. For this case, as $n \geq 4$ implies $k \geq 2$, each vertex in V_{k+1} has k distinct neighbours in V_{k+2}. Since all vertices outside V_{k+1} are fixed by induction, this guarantees V_{k+1} is fixed pointwise. $\square$

The above discussion suggests that finding double covers or antipodal quotients is a possible way to obtain new rigid flexible graphs. Note that $B_{n,n}$ is also a double cover of K_n. For any rigid flexible graph, can we always obtain another rigid flexible graph in this manner? Unfortunately, the answer is negative. For example, a double cover of $B_{5,5}$ is a graph with 20 vertices and of valency 4, and according to our data such a rigid flexible graph does not exist. The other way around, the antipodal quotient of a rigid flexible graph inherits the flexibility, but not necessarily the rigidity. We mentioned above that the quotient of Q_4 is non-rigid $K_{4,4}$. Another example is the quotient of the 1-skeleton of a dodecahedron, as shown in Figure 3.7. When we fix the neighbours x, y, z of the vertex o, swapping (x_1, x_2), (y_1, y_2), and (z_1, z_2) at the same time still gives a graph automorphism. In fact, the quotient is isomorphic to the Peterson graph. Despite these counter examples, we do have a general construction for the double cover of non-bipartite graphs. For this purpose, we need the following definition.

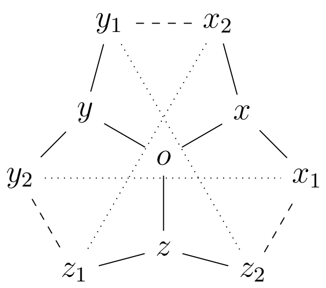

Figure 3.7: antipodal quotient of the dodecahedron graph

Definition 3.3. Suppose that Γ and Γ' are two simple graphs. The **direct product** of Γ and Γ', denoted by $\Gamma \times \Gamma'$, is a simple graph with vertex set the Cartesian product of vertex sets of Γ and Γ', and two vertices (v, v') and (u, u') are adjacent in $\Gamma \times \Gamma'$ if and only if v is adjacent to u in Γ and v' is adjacent to u' in Γ'.

With the above definition, it is easy to see that $K_2 \times K_n = B_{n,n}$, as illustrated in Figure 3.8 for the case of $n = 3$. Note that in the figure there is no edge between vertices with the same vertical or horizontal coordinate. Also note that the direct product of two edges is again two edges, laid out as a cross in the figure, which is part of the reason why graph theorists choose the symbol "$\times$" [15]. Therefore the direct product of two connected graphs is not necessarily connected. The following theorem about the connectedness of direct product is known as Weichsel's Theorem [15].

Theorem 3.4. Suppose that Γ and Γ' are two connected simple graphs with at least two vertices. If Γ and Γ' are both bipartite, then $\Gamma \times \Gamma'$ has exactly two components. If at least one of Γ and Γ' is not bipartite, then $\Gamma \times \Gamma'$ is connected.

Proof. The first part of the theorem is straightforward. For the second part, note that a simple graph is not bipartite if and only if there is an odd cycle in the graph. By exploiting such a cycle properly, the second part of the theorem follows. For a detailed proof, please refer to Theorem 5.9 in [15]. $\square$

Take Q_n^* as an example. Note that Q_n^* is bipartite if and only if n is even. For an even n, it is easy to see that $K_2 \times Q_n^*$ is the disjoint union of two copies of Q_n^*. For an odd n, it is also not hard to see that $K_2 \times Q_n^*$ recovers Q_n. The following proposition explains $K_2 \times Q_n^* = Q_n$ in a more general setting.

Proposition 3.5. A bipartite double cover of a simple connected non-bipartite graph Γ must be of the form $K_2 \times \Gamma$.

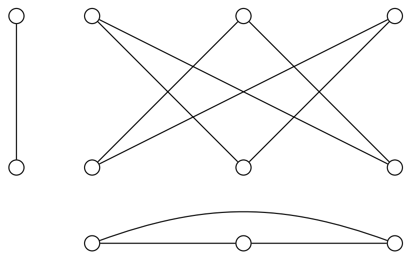

Figure 3.8: $K_2 \times K_3 = B_{3,3}$

Proof. Suppose that Γ' is a bipartite double cover of Γ. Since Γ is connected and non-bipartite, for any vertex v of Γ, there is a path of odd length starting and ending at v. Suppose that v_0 and v_1 are the two vertices in Γ' covering v. We can lift this path to a path of odd length in Γ' starting at v_0, ending at either v_0 or v_1. But this path can not end at v_0, otherwise there is an odd cycle in bipartite Γ'. The path of odd length from v_0 to v_1 shows that v_0 and v_1 must be in different partite sets of Γ'. This allows us to denote vertices of Γ' by (δ, v), where the first coordinate $\delta \in \{0, 1\}$ indicates the partite set it belongs to in Γ', and the second coordinate indicates the vertex v in Γ it covers.

Now consider the covering map from Γ' to Γ. If there is no edge between v and u in Γ, then in Γ' there is no edge between either $(0, v)$ and $(1, u)$, or $(1, v)$ and $(0, u)$. Suppose there is an edge between v and u. It must be covered by two edges in Γ', between either $(0, v)$ and $(1, u)$, or $(1, v)$ and $(0, u)$. Note that these two edges can not be double edges between two vertices, otherwise the covering map is not locally homeomorphic. Therefore, there is exactly one edge between $(0, v)$ and $(1, u)$, and exactly one between $(1, v)$ and $(0, u)$. The pattern of edges of Γ' shows that Γ' is isomorphic to $K_2 \times \Gamma$. $\square$

The examples of $K_2 \times K_n = B_{n,n}$ and $K_2 \times Q_n^* = Q_n$ suggest that the direct product of K_2 with a rigid flexible graph Γ could be rigid flexible. In fact this is essentially to show $\mathrm{Aut}(K_2 \times \Gamma) \cong C_2 \times \mathrm{Aut}(\Gamma)$, which does not hold for arbitrary graphs. When Γ is connected and bipartite, $K_2 \times \Gamma$ is two copies of Γ and hence has automorphism group $\mathrm{Aut}(\Gamma) \wr C_2$. Even for non-bipartite cases, if Γ has two vertices u and v with the same set of neighbours, then $(0, u)$ and $(0, v)$ have the same set of neighbours in $K_2 \times \Gamma$, and simply swapping $(0, u)$ and $(0, v)$ gives an automorphism of $K_2 \times \Gamma$ which is not in $C_2 \times \mathrm{Aut}(\Gamma)$. If we rule out these two types of obvious counterexamples, can we guarantee that $K_2 \times \Gamma$ is a rigid flexible graph? Unfortunately the answer is still negative.

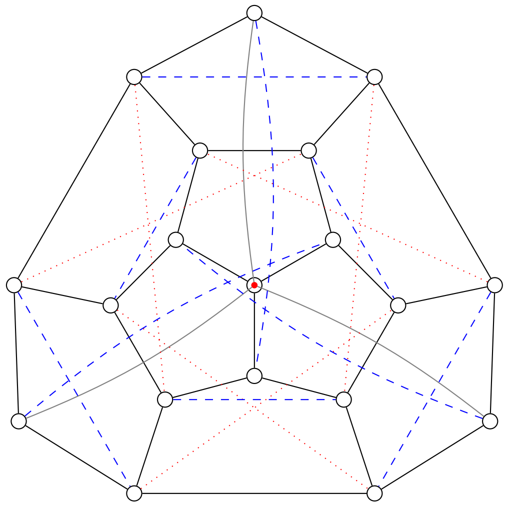

Figure 3.9: an exotic automorphism of K_2 $\times$ dodecahedron graph

Let Γ be the 1-skeleton of a dodecahedron, which is rigid flexible and non-bipartite. In the last part of Program 2, we verify that $|\operatorname{Aut}(K_2 \times \Gamma)| = 480$, while $|C_2 \times \operatorname{Aut}(\Gamma)| = 2 \cdot |\operatorname{Aut}(\Gamma)| = 240$. Hence we know that $K_2 \times \Gamma$ is not rigid, and it must have some exotic automorphisms, one of which is illustrated in Figure 3.9. In this particular drawing of Γ, the central red dot represents the antipodal vertex of the central vertex, and all the solid lines represent edges of Γ. For $K_2 \times \Gamma$, this figure should be viewed as there are two layers of vertices overlapping together, and each solid line represents two edges joining vertices at different layers. Let ρ be the permutation of vertices of $K_2 \times \Gamma$ such that ρ swaps any two vertices in the first layer joined by a blue dashed line, any two vertices in the second layer joined by a red dotted line, and fixes all other vertices. The central red dot denotes that two central vertices are joined by a red dotted line. Note that if we pick two vertices joined by a solid line, by following lines with different colours, we always obtain two vertices joined by a solid line. This is essentially saying ρ preserves edges in $K_2 \times \Gamma$. It is routine to verify that ρ preserves non-edges as well, and hence ρ is an automorphism of $K_2 \times \Gamma$. The different actions of ρ at different layers show that $\rho \notin C_2 \times \operatorname{Aut}(\Gamma)$, and judging from the cardinality we know that ρ and $C_2 \times \operatorname{Aut}(\Gamma)$ generate $\operatorname{Aut}(K_2 \times \Gamma)$.

Chapter 4

Classification Theorem

In this chapter we generalize the classification theorem of CAT(0) platonic polygonal complexes in [20]. Before we describe the main theorem, let us first introduce the main machinery, the theory of triangles of groups. Throughout the chapter, X is a platonic polygonal complex, G is the automorphism group of X, and fev is a flag in X. Note that the stabilizers G_{fev}, G_{fe}, G_{fv}, G_{ev}, G_f, G_e, G_v form a commutative diagram as shown in Figure 4.1, where each arrow is an injective group homomorphism. Stallings calls a diagram like this a **triangle of groups** in [26].

With G and these stabilizers, we can recover the complex X as follows. First we construct a complex with left cosets $G/G_f \coprod G/G_e \coprod G/G_v$ as vertices, $G/G_{fe} \coprod G/G_{fv} \coprod G/G_{ev}$ as edges, and G/G_{fev} as triangles, where a vertex, an edge, and a triangle are incident if one contains another as a set. The complex we build is actually the barycentric subdivision of X. By gluing triangles at vertices of type G/G_f, we obtain the original complex

$$\begin{tikzcd}
 & & G_f & & \\
 & G_{fe} \arrow[ur] \arrow[ddl] & & G_{fv} \arrow[ul] \arrow[ddr] & \\
 & & G_{fev} \arrow[uu] \arrow[ul] \arrow[ur] \arrow[dll] \arrow[d] \arrow[drr] & & \\
G_e & & G_{ev} \arrow[ll] \arrow[rr] & & G_v
\end{tikzcd}$$

Figure 4.1: a triangle of groups

X. This suggests that we could study a platonic polygonal complex through its triangle of stabilizers. Then the first question would be, if we start with a triangle of groups, how could we obtain the group G?

When X is simply-connected, the theory of covering spaces (see [17]) shows that G is the colimit (see [23] for the definition) of its triangle of stabilizers, as there is an one-to-one correspondence between quotients of G and spaces covered by X. Unfortunately, the colimit of an arbitrary triangle of groups does not always behave nicely. Here is a well-known example: G_{fev} is the trivial group. G_{fe}, G_{fv}, and G_{ev} are infinite cyclic groups generated by a, b, and c respectively. And the other three groups are defined as:

$$G_f = \langle\, a, b \mid aba^{-1} = b^2 \,\rangle$$
$$G_e = \langle\, c, a \mid cac^{-1} = a^2 \,\rangle$$
$$G_v = \langle\, b, c \mid bcb^{-1} = c^2 \,\rangle$$

These groups form a triangle of groups, while the colimit group G is the group generated by a, b, and c under three relators in G_f , G_e, and G_v. Although not immediately obvious, this sibling of the Higman group is actually a trivial group.

For the moment, forget about the geometric meaning of these stabilizers, and let G be the colimit of the triangle of groups in Figure 4.1. A triangle of groups is called **developable** if the natural maps $G_f \to G$, $G_e \to G$, and $G_v \to G$ are injective. As the above example suggests, a triangle of groups is not always developable. Stallings gives a sufficient condition for a triangle of groups to be developable in [26]. First he defines the angle between two subgroups in a group. Let A, B, and $C \leq A \cap B$ be subgroups of G, and consider the natural homomorphism ϕ from the amalgamated free product $A *_C B$ to G. We can define the length of $g \in A *_C B$ as the smallest m such that $g = c_1 c_2 \dots c_m$, where each $c_i \in A \cup B$. Suppose that $g = a_1 b_1 a_2 b_2 \dots$ is a shortest nontrivial element in $\mathrm{Ker}(\phi)$, where $a_i \in A$ and $b_j \in B$. If the last letter in the product is a_n, then $a_n g a_n^{-1}$ is again in $\mathrm{Ker}(\phi)$ and has shorter length. Therefore the length of g must be an even number $2n$. We define the **angle** between A and B with respect to C, denoted by $(A, B; C)$, to be π/n. In case $\mathrm{Ker}(\phi)$ is trivial, we define $(A, B; C)$ to be 0.

Take the dihedral group $D_n = \langle\, a, b \mid a^2 = b^2 = (ab)^n = 1 \,\rangle$ as an example. We can think of D_n as symmetries of regular n-gon, a as a reflection fixing a vertex v, and b as a reflection fixing an edge containing v. Then the angle between two reflection axes is π/n. Let A and B be the subgroups generated by a and b respectively. Consider the homomorphism from $A * B$ to D_n. The shortest nontrivial element in the kernel is $(ab)^n$,

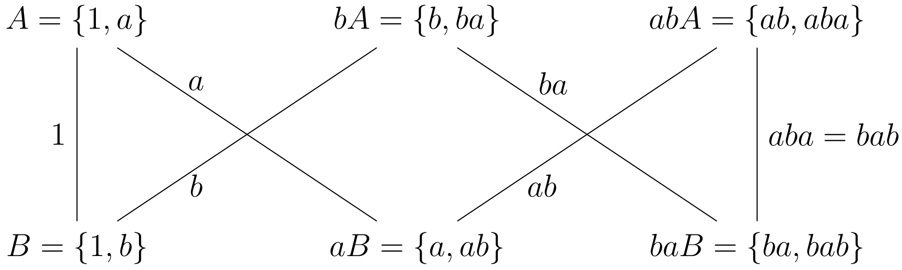

Figure 4.2: $\Gamma(A, B; \{1\})$ of D_3

of which the length is $2n$. Therefore, the angle $(A, B; \{1\})$ is π/n, the same as the angle between reflection axes of a and b. This to some extent justifies the term "angle".

In general, it is quite hard to calculate the angle. In some occasions, the following combinatorial method makes calculation easier. Suppose A, B, and $C \leq A \cap B$ are subgroups of G. We define a bipartite graph $\Gamma(A, B; C)$, where the vertex set consists of left cosets G/A and G/B, and the edge set consists of left cosets G/C. An edge gC in $\Gamma(A, B; C)$ joins gA and gB. For example, $\Gamma(A, B; \{1\})$ of the dihedral group D_3 is illustrated in Figure 4.2. Note that the girth of $\Gamma(A, B; C)$ is exactly the length of a shortest nontrivial element in the kernel of $A *_C B \to G$. Since $\Gamma(A, B; C)$ is a bipartite graph, this also explains why the shortest length is always even.

Now we look at the triangle of groups in Figure 4.1. In G_f, we can calculate at the angle $(G_{fe}, G_{fv}; G_{fev})$. For convenience, we call it the angle at G_f. Similarly, we can calculate the angles at G_e and G_v. A triangle of group is called **non-spherical** if the sum of these three angles is less than or equal to π. Stallings gives the following theorem.

Theorem 4.1. [26] Any non-spherical triangle of groups is developable.

Let us examine the angle sum of the sibling of the Higman group on the previous page. For the angle at G_f, look at the homomorphism $G_{fe} *_{G_{fev}} G_{fv} \to G_f$. Note that $G_{fe} *_{G_{fev}} G_{fv}$ is the free group generated by a and b, and the element $aba^{-1}b^{-2}$ of length 4 is one of the shortest nontrivial elements in the kernel. Therefore the angle at G_f is $\pi/2$. Similarly the angles at G_e and G_v are both $\pi/2$, and the sum of angles is greater than π. Hence this triangle of groups is not non-spherical, or spherical we should say.

While the proof of this theorem is not in the scope of this thesis, the example of a triangle group (l, m, n) might help to give some insight. A triangle group (l, m, n) can be presented as $\langle\, a, b, c \mid a^2 = b^2 = c^2 = (ab)^l = (bc)^n = (ca)^m = 1 \,\rangle$. Let G_{fev} be the trivial group. G_{fe}, G_{fv}, and G_{ev} are C_2 generated by a, b, and c respectively. G_f is the

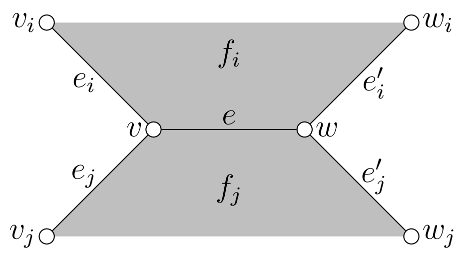

Figure 4.3: local bijection

dihedral group generated by a and b, G_e is the dihedral group generated by c and a, and G_v is the dihedral group generated by b and c. These groups form a triangle of groups with sum of angles $\pi/l + \pi/m + \pi/n$, and the colimit of this triangle of groups is exactly the triangle group (l, m, n). While the sum of angles is π or less than π, the triangle group (l, m, n) gives a tessellation of Euclidean plane or hyperbolic plane respectively. This suggests that a non-spherical triangle of groups indeed carries certain structure of non-positive curvature, and this keeps the colimit group from collapsing.

Now we go back to the discussion of complexes. Our plan is to classify platonic polygonal complexes through classifying possible triangles of stabilizers, with respect to a symmetric rigid link L. By Proposition 2.9, the rigidity of L implies the rigidity of complexes. Therefore $g \in G_v$ is determined by its action on the edges incident to v, which correspond to the vertices in L. Then G_v induces a flag-transitive action on L, where a flag in a graph is a pair of incident vertex and edge. While L discloses plenty of information about G_v, what about other stabilizers? This leads to the next topic.

Suppose that v and w are two vertices joined by an edge e in a platonic polygonal complex X with link L, a symmetric rigid graph of valency m with n vertices. The edge valency of X is the same as the valency of L, so let $f_1, \ldots, f_m$ be the m faces incident to e. Each face f_i has consecutive vertices v_i, v, w, w_i joined by e_i, e, e'_i as illustrated in Figure 4.3. This partially defines a function $\phi_{v,w}$ from edges incident to v to edges incident to w, where $\phi_{v,w}(e_i) = e'_i$ and $\phi_{v,w}(e) = e$. We can identify the edge e as a vertex in the link of either v or w. Then $\phi_{v,w}$ can be thought of as a function from the neighbourhood of e in $L(X, v)$ to the neighbourhood of e in $L(X, w)$. We hope that $\phi_{v,w}$ can be extended to a graph isomorphism from $L(X, v)$ to $L(X, w)$. If $\phi_{v,w}$ is extendable, then $\phi_{v,w}$ gives a bijection between edges incident to v and edges incident to w, and a bijection between neighbours of v and neighbours of w. We can alternatively define $\phi_{v,w}$ as a bijection between neighbours of v and w. In particular, we have $\phi_{v,w}(v_i) = w_i$ and $\phi_{v,w}(w) = v$.

Definition 4.2. A complex X is **good** if it satisfies the following:

(1) X is a simply-connected platonic polygonal complex.

(2) The link of X is a symmetric rigid graph.

(3) For two adjacent vertices v and w, the function $\phi_{v,w}$ is extendable.

Note that condition (3) does not depend on the choice of adjacent vertices, so the goodness of a complex is well-defined. Moreover, if $\phi_{v,w}$ is extendable, the rigidity in condition (2) implies that $\phi_{v,w}$ is uniquely extendable. For a platonic polygonal complex, the function $\phi_{v,w}$ is not always extendable. We do not have many clues for these non-extendable cases yet. If we assume the link is flexible, then every permutation of neighbours of a vertex in the link extends to a graph automorphism, which implies $\phi_{v,w}$ is extendable. When $\phi_{v,w}$ is extendable, it interacts nicely with complex automorphisms.

Proposition 4.3. Suppose X is a good complex, and v and w are two adjacent vertices in X. Then for any automorphism g of X, we have

$$g \circ \phi_{v,w} = \phi_{g(v),g(w)} \circ g.$$

In other words, any automorphism of X preserves the bijection.

Proof. Suppose f_i is a face with consecutive vertices v_i, v, w, w_i. By definition $\phi_{v,w}(v_i) = w_i$. If we apply g to this face, then we have a face with consecutive vertices $g(v_i)$, $g(v)$, $g(w)$, $g(w_i)$, and therefore $\phi_{g(v),g(w)}(g(v_i)) = g(w_i) = g(\phi_{v,w}(v_i))$. Also we have $\phi_{g(v),g(w)}(g(w)) = g(v) = g(\phi_{v,w}(w))$. Note $g \circ \phi_{v,w}$ and $\phi_{g(v),g(w)} \circ g$ both define a graph isomorphism from $L(X, v)$ to $L(X, g(w))$, and these two isomorphisms coincide at vertex vw and all its neighbours like vv_i. By the rigidity of L, we have $g \circ \phi_{v,w} = \phi_{g(v),g(w)} \circ g$. $\square$

Now we study how these bijections affect the local structure around a face. Suppose f is a d-gonal face in X with vertices $v_0, v_1, \ldots, v_{d-1}$ listed in order, and N_i is the set of neighbours of v_i. Let v be $v_0 = v_d$, e be the edge v_0v_1, and $\phi_i = \phi_{v_i,v_{i+1}}$ for $0 \leq i \leq d-1$. We will stick to this notation in the following few paragraphs.

Definition 4.4. The **holonomy** at v along e and f is defined as $\Phi = \phi_{d-1} \circ \phi_{d-2} \circ \cdots \circ \phi_1 \circ \phi_0$. Φ is an automorphism of $L(X, v)$, as well as a permutation of N_0.

Note when d is even, Φ fixes v_1 and v_{d-1}, and when d is odd, Φ swaps v_1 and v_{d-1}. The holonomy element is not necessarily the identity, nor an arbitrary permutation. The constraint of holonomy will help us to classify possible triangles of stabilizers. For the rest of the chapter, Φ denotes the holonomy at v along e and f.

Proposition 4.5. Suppose that X is a good d-gonal complex with a flag fev, and Φ is the holonomy at v along e and f. Let $g \in G_{fev}$, $r \in G_{fv} - G_{fev}$, $s_0 \in G_{fe} - G_{fev}$, and $s = s_0 \circ \phi_0$. Then we have the following relations:

(1) $g\Phi g^{-1} = \Phi$

(2) $r\Phi r^{-1} = \Phi^{-1}$

(3) $s\Phi s^{-1} = \Phi^{-1}$

(4) $(rs_0)^d\Phi = (rs)^d$

Proof. For (1), note that $g \in G_{fev}$ also fixes all other vertices and edges in f. We have $g \circ \phi_{v_i,v_{i+1}} = \phi_{g(v_i),g(v_{i+1})} \circ g = \phi_{v_i,v_{i+1}} \circ g$, and therefore $g \circ \phi_i = \phi_i \circ g$ for all i. Hence $g \circ \Phi = g \circ \phi_{d-1} \circ \phi_{d-2} \circ \cdots \circ \phi_1 \circ \phi_0 = \phi_{d-1} \circ \phi_{d-2} \circ \cdots \circ \phi_1 \circ \phi_0 \circ g = \Phi \circ g$. We can rewrite the equation as $g \circ \Phi \circ g^{-1} = \Phi$. In other words, Φ centralizes G_{fev}.

For (2), note that $r \in G_{fv} - G_{fev}$ acts as a reflection in f, fixing v_0, and swapping v_i with v_{d-i}. We have $r \circ \phi_{v_i,v_{i+1}} = \phi_{r(v_i),r(v_{i+1})} \circ r = \phi_{v_{d-i},v_{d-i-1}} \circ r$, and therefore $r \circ \phi_i = \phi_{d-i-1}^{-1} \circ r$ for all i. Hence $r \circ \Phi = r \circ \phi_{d-1} \circ \phi_{d-2} \circ \cdots \circ \phi_1 \circ \phi_0 = \phi_0^{-1} \circ \phi_1^{-1} \circ \cdots \circ \phi_{d-2}^{-1} \circ \phi_{d-1}^{-1} \circ r = \Phi^{-1} \circ r$. We can rewrite the equation as $r \circ \Phi \circ r^{-1} = \Phi^{-1}$.

For (3), note that $s_0 \in G_{fe} - G_{fev}$ acts as a reflection in f, fixing e, and swapping v_i with v_{d-i+1}, in particular v_0 and v_1. We have $s_0 \circ \phi_{v_i,v_{i+1}} = \phi_{s_0(v_i),s_0(v_{i+1})} \circ s_0 = \phi_{v_{d-i+1},v_{d-i}} \circ s_0$, and therefore $s_0 \circ \phi_i = \phi_{d-i}^{-1} \circ s_0$ for all i. Note that $s = s_0\phi_0$ induces a permutation of N_0. Consider $s\Phi = s_0 \circ \phi_0 \circ \phi_{d-1} \circ \phi_{d-2} \circ \cdots \circ \phi_1 \circ \phi_0 = \phi_0^{-1} \circ \phi_1^{-1} \circ \cdots \circ \phi_{d-2}^{-1} \circ \phi_{d-1}^{-1} \circ s_0 \circ \phi_0 = \Phi^{-1} \circ s$. Rewrite the equation and then we have $s \circ \Phi \circ s^{-1} = \Phi^{-1}$.

For (4), note that rs_0 acts as a rotation in f, which rotates v_{i+1} to v_i. We have $rs_0 \circ \phi_{v_i,v_{i+1}} = \phi_{rs_0(v_i),rs_0(v_{i+1})} \circ rs_0 = \phi_{v_{i-1},v_i} \circ rs_0$, and therefore $rs_0 \circ \phi_i = \phi_{i-1} \circ rs_0$. Now consider the map $(rs_0)^k \circ \phi_{k-1} \circ \phi_{k-2} \circ \cdots \circ \phi_0$. The ϕ_i's part of this map sends N_0 to N_k, while $(rs_0)^k$ rotates N_k back to N_0. Therefore this map is a permutation of N_0, and actually can be rewritten as $(rs)^k$. We prove this by induction. The base case $rs_0 \circ \phi_0 = rs$ comes from the definition of s. Assume that $(rs_0)^i \circ \phi_{i-1} \circ \phi_{i-2} \circ \cdots \circ \phi_0 = (rs)^i$. Then $(rs_0)^{i+1} \circ \phi_i \circ \phi_{i-1} \circ \cdots \circ \phi_0 = (rs_0)((rs_0)^i \circ \phi_i) \circ \phi_{i-1} \circ \cdots \circ \phi_0 = (rs_0)((rs_0)^{i-1} \circ \phi_{i-1} \circ (rs_0)) \circ \phi_{i-1} \circ \cdots \circ \phi_0 = \cdots = (rs_0)(\phi_0 \circ (rs_0)^i) \circ \phi_{i-1} \circ \cdots \circ \phi_0 = (rs_0 \circ \phi_0) \circ (rs_0)^i \circ \phi_{i-1} \circ \phi_{i-2} \circ \cdots \circ \phi_0 = (rs) \circ (rs)^i = (rs)^{i+1}$. By induction, we know the claim is true. In particular, if we take $k = d$, then we have $(rs_0)^d \circ \Phi = (rs_0)^d \circ \phi_{d-1} \circ \phi_{d-2} \circ \cdots \circ \phi_1 \circ \phi_0 = (rs)^d$. $\square$

We are ready to construct a triangle of stabilizers of a d-gonal good complex with respect to a given symmetric rigid link L of n vertices. To start with, we should choose

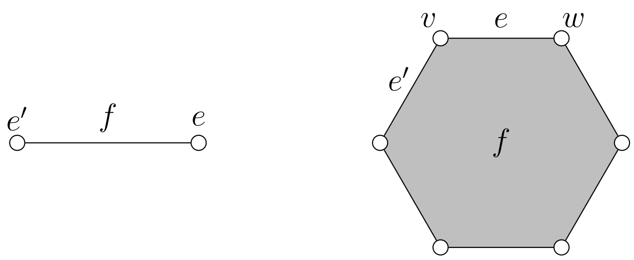

Figure 4.4: link versus complex

G_v to be a flag-transitive subgroup of $\mathrm{Aut}(L)$, as G acts transitively on fev flags. Next we specify a pair of adjacent vertex e and e' in L, where f is the edge joining them. G_{ev}, G_{fv}, and G_{fev} are then defined respectively as the stabilizer in G_v of e, the unordered pair $\{e, e'\}$, and the ordered pair (e, e'), as illustrated in Figure 4.4. Introducing e' seems rather redundant than just using f, but it is actually how we code the elements of G_v in programming, which is a permutation of n vertices of L.

Now we look at G_e. Suppose v and w are the two vertices incident to e in the complex. The rigidity of the complex shows that $g \in G_e$ is completely determined by its action on v and its neighbours. Moreover, the bijection map $\phi_{v,w}$ between neighbours of v and w makes it clear that G_e is a subgroup of $C_2 \times S_{n-1}$, where $(\delta, \rho) \in C_2 \times S_{n-1}$ sends an incident edge e_* of v to $\phi_{v,w}^{\delta} \circ \rho(e_*)$. Note that $G_e \cap (\{0\} \times S_{n-1})$ should be identified as G_{ev}. Therefore, we choose G_e to be an index 2 supergroup over $\{0\} \times G_{ev}$ in $C_2 \times S_{n-1}$ such that $G_e \cap (\{0\} \times S_{n-1}) = \{0\} \times G_{ev}$. Then G_{fe} is defined as $\{(\delta, \rho) \in G_e \mid \rho \text{ fixes } e'\}$. The above notation is slightly complicated. To simplify it, consider the projection map $\pi : C_2 \times S_{n-1} \to S_{n-1}$, and let $\overline{G_e} = \pi(G_e)$. If $(1, 0) \in G_e$, then G_e is essentially $C_2 \times G_{ev}$ and $\overline{G_e} = G_{ev}$. If $(1, 0) \notin G_e$, then π gives an isomorphism between G_e and $\overline{G_e}$. Therefore, given any index 1 or 2 supergroup of G_{ev} in S_{n-1}, namely $\overline{G_e}$, we have a corresponding G_e. When the index is 1, $G_e = C_2 \times G_{ev}$, and $G_{fe} = C_2 \times G_{fev}$. When the index is 2, $G_e = \overline{G_e}$, and $G_{fe} = \{g \in \overline{G_e} \mid g \text{ fixes } e'\}$. Note that for every $s_0 \in G_{fe} - G_{fev}$, s_0 is of the form $(1, \rho) \in C_2 \times S_{n-1}$, and $s = s_0 \circ \phi_0 = (1, \rho) \circ (1, id) = (0, \rho) = \pi(s_0)$.

For example, let L be the complete graph K_4 with vertex set $\{e, e', e_1, e_2\}$, and $G_v = A_4$ acts flag-transitively on K_4. Then $G_{ev} = (G_v)_e = \langle\, (e', e_1, e_2)\, \rangle$, $G_{fv} = (G_v)_{\{e,e'\}} = \langle\, (e, e')(e_1, e_2)\, \rangle$, and $G_{fev} = (G_v)_{(e,e')}$ is trivial. Let us consider the index 2 supergroup $\overline{G_e}$ over G_{ev} in S_{4-1}, the symmetric group on $\{e', e_1, e_2\}$. Then $G_e = \overline{G_e}$ must be S_3, and $G_{fe} = (S_3)_{e'} = \langle\, (e_1, e_2)\, \rangle$. Note that $(e_1, e_2) \in G_e - G_{ev}$, and its actual action is sending e_1 to $\phi_{v,w}(e_2)$, e_2 to $\phi_{v,w}(e_1)$, and e' to $\phi_{v,w}(e')$, which indicates that f is stabilized.

We have defined every group in the triangle of stabilizers except G_f. From the diagram, we know G_f should be a quotient of the amalgamated free product $G_{fv} *_{G_{fev}} G_{fe}$. Note that $(G_{fv} *_{G_{fev}} G_{fe})/G_{fev} \cong C_2 * C_2$, as G_{fev} is an index 2 normal subgroup in both G_{fv} and G_{fe}. Moreover, G_f/G_{fev} is the dihedral group of order $2d$, since G_f induces the symmetry group of the d-gonal face f with kernel G_{fev}. This suggests that for every $r \in G_{fv} - G_{fev}$ and $s_0 \in G_{fe} - G_{fev}$, the element $(rs_0)^d$ should be identified with some element in G_{fev}. This is where the holonomy Φ kicks in.

We should define G_f in a way that a holonomy element exists and satisfies all the relations in Proposition 4.5. Choose an arbitrary $r \in G_{fv} - G_{fev}$ and $s_0 \in G_{fe} - G_{fev}$. Previously s is defined as $s_0 \circ \phi_0$, which is exactly $\pi(s_0)$. Hence we define $s = \pi(s_0)$. Suppose that there exists $\Phi \in \mathrm{Aut}(L)$ satisfying the following:

(1) $g\Phi g^{-1} = \Phi, \forall g \in G_{fev}$

(2) $r\Phi r^{-1} = \Phi^{-1}$

(3) $s\Phi s^{-1} = \Phi^{-1}$

(4) $(rs)^d\Phi^{-1} \in G_{fev}$

As suggested by Proposition 4.5, we would like to identify $(rs_0)^d$ with $(rs)^d\Phi^{-1}$. We also mentioned that $(rs_0)^d$ should be identified with an element in G_{fev}. With condition (4), now it makes sense to define

$$G_f := (G_{fv} *_{G_{fev}} G_{fe}) \,/\, \langle\langle\, (rs_0)^d = (rs)^d\Phi^{-1} \,\rangle\rangle.$$

Note that r swaps e and e', s fixes e and e', and G_{fev} fixes e and e'. Condition (4) implies that Φ fixes e and e' for even d, and swaps e and e' for odd d, a property we expect a holonomy to have. In fact if $G_v = \mathrm{Aut}(L)$, this property is equivalent to condition (4). For $G_v \lneq \mathrm{Aut}(L)$, this property does not guarantee condition (4).

Proposition 4.6. Suppose that r, s_0, s, Φ, G_f are defined as above. Then for $r' \in G_{fv} - G_{fev}$, $s_0' \in G_{fe} - G_{fev}$, and $s' = \pi(s_0')$, we have

(1) G_f/G_{fev} is the dihedral group of order $2d$.

(2) $r'\Phi r'^{-1} = \Phi^{-1}$

(3) $s'\Phi s'^{-1} = \Phi^{-1}$

(4) $(r's_0')^d = (r's')^d\Phi^{-1}$ in G_f.

In particular, the definition of G_f is independent of the choices of r and s_0.

Proof. (1) is clear from the definition of G_f. For (2), (3), (4), note that $r^{-1}r'$ and $s_0^{-1}s_0'$ are both in G_{fev}. Therefore $\exists g, h \in G_{fev}$ such that $r' = rg$ and $s_0' = s_0 h$. Also note that $s' = \pi(s_0') = \pi(s_0 h) = \pi(s_0)h = sh$. Thus we have

$$r'\Phi r'^{-1} = rg\Phi g^{-1}r^{-1} = r\Phi r^{-1} = \Phi^{-1},$$

$$s'\Phi s'^{-1} = sh\Phi h^{-1}s^{-1} = s\Phi s^{-1} = \Phi^{-1}.$$

To show (4), note that it suffices to show the following:

$$(r's_0')^{-d}(rs_0)^d = \Phi(r's')^{-d}(rs)^d\Phi^{-1}$$

By (2) and (3) of Φ's property and (2) and (3) of this proposition, we can move Φ to the right of $(r's')^{-d}(rs)^d$ and still get Φ after swapping between Φ^{-1} and Φ $2d$ times. Therefore the above equation is equivalent to

$$(rgs_0h)^{-d}(rs_0)^d = (rgsh)^{-d}(rs)^d.$$

Note that for any $x \in G_{fev}$ we have $s_0^{-1}xs_0 = s^{-1}xs$. We can expend $(rgs_0h)^{-d}(rs_0)^d$ and use the above relation to replace s_0 to s, starting from the middle of the product. Eventually we can rewrite $(rgs_0h)^{-d}(rs_0)^d$ as $(rgsh)^{-d}(rs)^d$. $\square$

Suppose that X is a good complex, and G is the automorphism group of X. If G has a proper subgroup G' which also acts flag-transitively on X, then the stabilizers of G' also form a triangle of groups, and the complex built by using this triangle of groups also recovers the original complex. Conversely, if we have two triangles of groups with triples $(G_v, \overline{G_e}, \Phi)$ and $(H_v, \overline{H_e}, \Phi)$ where $G_v \leq H_v$ and $\overline{G_e} \leq \overline{H_e}$, then these two triples determine the same polygonal complex. We define a partial ordering on these triples by $(G_v, \overline{G_e}, \Phi) \leq (H_v, \overline{H_e}, \Phi')$ if $G_v \leq H_v$, $\overline{G_e} \leq \overline{H_e}$, and $\Phi = \Phi'$. For classification purpose, we only care about maximal $(G_v, \overline{G_e}, \Phi)$. Also note that an automorphism of L acts on triples $(G_v, \overline{G_e}, \Phi)$ by conjugation. Any two conjugate triples also determine the same complex. Now we are ready to state the classification theorem.

Theorem 4.7. Let L be a symmetric rigid graph of n vertices with girth l. If $\frac{1}{d} + \frac{1}{l} \leq \frac{1}{2}$, then there is a bijection between good d-gonal complexes with link L and conjugacy classes of maximal triples $(G_v, \overline{G_e}, \Phi)$, where G_v is a flag-transitive subgroup of $\mathrm{Aut}(L)$, $\overline{G_e}$ is an index 1 or 2 supergroup over G_{ev} in S_{n-1}, and $\Phi \in \mathrm{Aut}(L)$ satisfies the following:

(1) $g\Phi g^{-1} = \Phi$, $\forall g \in G_{fev}$

(2) $r\Phi r^{-1} = \Phi^{-1}$, for some $r \in G_{fv} - G_{fev}$

(3) $s\Phi s^{-1} = \Phi^{-1}$, where s is the image of some $s_0 \in G_{fe} - G_{fev}$ under the projection map $\pi : C_2 \times S_{n-1} \to S_{n-1}$

(4) $(rs)^d \Phi^{-1} \in G_{fev}$

Proof. We know that each good d-gonal complex with link L corresponds to a conjugacy class of maximal triples $(G_v, \overline{G_e}, \Phi)$. Conversely, by Proposition 4.6, each maximal triple $(G_v, \overline{G_e}, \Phi)$ determines a triangle of groups, independent of the choices of r and s_0. Let us look at the angles of the resulting triangle of groups. Since G_f/G_{fev} is the dihedral group of order $2d$, the graph $\Gamma(G_{fv}, G_{fe}; G_{fev})$ is a cycle of length $2d$. Therefore the angle at G_f is $\frac{\pi}{d}$. Now we look at $\Gamma(G_{fv}, G_{ev}; G_{fev})$ in G_v. Since G_v acts flag-transitively on L, $\Gamma(G_{fv}, G_{ev}; G_{fev})$ is the same as the barycentric subdivision of L. The girth of this graph doubles the girth of L. Therefore the angle at G_v is $\frac{\pi}{l}$, where l is the girth of L. As for $\Gamma(G_{fe}, G_{ev}; G_{fev})$, it is not hard to see that it is a complete bipartite graph $K_{2,m}$, where m is the valency of L. The girth of $K_{2,m}$ is 4, and therefore the angle at G_e is $\frac{\pi}{2}$. By Theorem 4.1, if the sum of angles $\frac{\pi}{d} + \frac{\pi}{l} + \frac{\pi}{2} \leq \pi$, or equivalently $\frac{1}{d} + \frac{1}{l} \leq \frac{1}{2}$, then these stabilizers inject into the colimit group G, and their cosets form a simplicial complex. The above discussion also shows that this complex has the right links, so we can glue the faces at G/G_f to get a good d-gonal complex with link L. Note that conjugate maximal triples result in isomorphic complexes, and this completes the proof. □

In addition to platonic polygonal complexes, we are also interested in complexes with the following properties: the automorphism group has two orbits on each of fev, fe, fv, f, and acts transitively on every other type of partial flag. We call such a complex a **half platonic complex**. Take an octahedron as an example. We can insert three bisecting squares to the octahedron, as illustrated in Figure 4.5. The resulting complex is still highly symmetric, and indeed a half platonic complex.

In [29], Valle investigates platonic complexes with octahedral graphs as links. An octahedral graph is a graph where every vertex is connected to every other vertex except one, as illustrated in Figure 4.5. In such a complex, we can find a cycle such that no consecutive three vertices form a corner of a face. These cycles are called holes, and the automorphism group acts transitively on these holes. If we attach a face to each hole, then we can also obtain a half platonic complex, as in the octahedron example. Therefore Valle's work provides lots of examples of half platonic complexes.

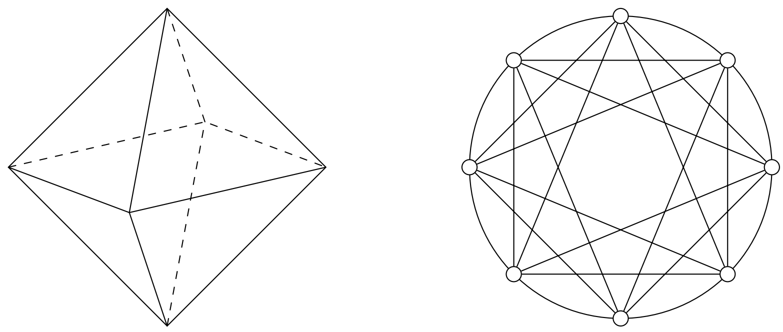

Figure 4.5: an octahedron and an octahedral graph

$$
\begin{tikzcd}
 & & & & G_v & & & & \\
 & & G_{f'v} \arrow[urr] \arrow[ddll] & & G_{ev} \arrow[u] \arrow[dd] & & G_{fv} \arrow[ull] \arrow[ddrr] & & \\
 & & & G_{f'ev} \arrow[uur] \arrow[ul] \arrow[ur] \arrow[dlll] \arrow[dl] \arrow[dr] & & G_{fev} \arrow[uul] \arrow[ul] \arrow[ur] \arrow[dl] \arrow[dr] \arrow[drrr] & & & \\
G_{f'} & & G_{f'e} \arrow[ll] \arrow[rr] & & G_e & & G_{fe} \arrow[ll] \arrow[rr] & & G_f
\end{tikzcd}
$$

Figure 4.6: double triangles of groups

We would like to classify half platonic complexes by imitating the method in this chapter. Now the stabilizers of partial flags form a diagram like Figure 4.6. Under certain curvature conditions we can guarantee the developability of the diagram. And since we have two different types of faces, there should be two holonomies involved in the construction. This project is still ongoing.

Chapter 5

Tensor Product

While the theory of triangles of groups provides a computational tool to find platonic complexes, it does not give much clue about how to explicitly construct examples, due to the difficulty of understanding the colimit group, especially for non-CAT(0) situations. In [20], standard examples of platonic complexes with K_n links and d-gonal faces are constructed as follows. By assuming the holonomy Φ to be trivial, $G_v = S_n$, $G_e = S_{n-1} \times S_2$, and $G_f = S_{n-2} \times D_d$, the colimit of this triangle of groups is a Coxeter group of n generators with the following Coxeter diagram,

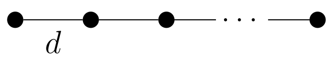

where G_v, G_e, and G_f can be obtained by removing the first, second, and third generator from the left respectively. This sort of construction does not seem to work in general for other rigid flexible links such as $B_{n,n}$ and Q_n. Note that $B_{n,n}$ can be obtained from the direct product of K_2 with K_n, and Q_n can be obtained from the Cartesian product (defined on p. 49) of n copies of K_2. In the following two chapters, we develop two types of complex products, which interact nicely with the above two products of link graphs. This allows us to construct platonic complexes with $B_{n,n}$ and Q_n links.

Suppose that $\bullet$ is certain type of graph product such that $V(\Gamma \bullet \Gamma') = V(\Gamma) \times V(\Gamma')$, and we want to define a complex product $*$ which interacts with $\bullet$ nicely. More specifically, we would like $*$ to have the following property: for any complexes X and X', and for any vertices $v \in X$ and $v' \in X'$, we have

$$L(X, v) \bullet L(X', v') \cong L(X * X', (v, v')).$$

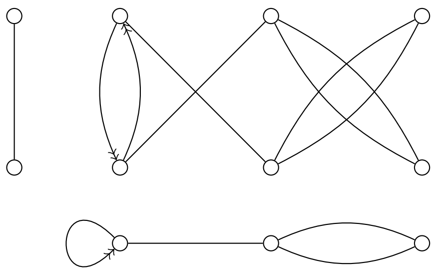

Figure 5.1: tensor product of non-simple graphs

Here we have already assumed that $V(X * X') = V(X) \times V(X')$. The above property provides sufficient information about how the complex product $*$ shall be defined. If we assume the 1-skeletons of X and X' are simple graphs, by considering the vertex sets of two link graphs in the equation, we have

$$\{\text{neighbours of } v \text{ in } X\} \times \{\text{neighbours of } v' \text{ in } X'\} = \{\text{neighbours of } (v, v') \text{ in } X * X'\},$$

which can be interpreted as two vertices (v, v') and (u, u') are adjacent in $X * X'$ if and only if v is adjacent to u in X and v' is adjacent to u' in X'. This is exactly the condition in Definition 3.3, which means the 1-skeleton of $X * X'$ should be the direct product of 1-skeletons of X and X'. Since the 1-skeletons of complexes are not necessarily simple, we shall generalize the direct product to suit arbitrary graphs.

Definition 5.1. Suppose that Γ and Γ' are two arbitrary graphs with edge sets $E(\Gamma) = \{e_\alpha \mid \alpha \in A\}$ and $E(\Gamma') = \{e_\beta \mid \beta \in B\}$. The **tensor product** of Γ and Γ', denoted by $\Gamma \otimes \Gamma'$, is a graph with vertex set $V(\Gamma \otimes \Gamma') = V(\Gamma) \times V(\Gamma')$, and edge set

$$E(\Gamma \otimes \Gamma') = \left\{ e^\delta_{\alpha,\beta} \mid \alpha \in A, \beta \in B, \delta \in \{0, 1\} \right\},$$

where $e^\delta_{\alpha,\beta}$ is an edge joining (v_0, v'_δ) and $(v_1, v'_{1-\delta})$, given e_α joins v_0 and v_1 in Γ, and e_β joins v'_0 and v'_1 in Γ'.

Note that for simple graphs, the tensor product defined above is exactly the direct product of graphs. Like direct product, each pair of edges from two factors generates two edges in the tensor product, even when loops are involved, as illustrated in Figures 5.1 and 5.2. In some literatures such as [15], direct product is defined over graphs without parallel edges but admitting loops. In such definition, a loop serves as the identity of direct product. In particular a loop times an edge is an edge, and a loop times a loop is again a loop, while in our definition a loop times an edge is two parallel edges, and a loop

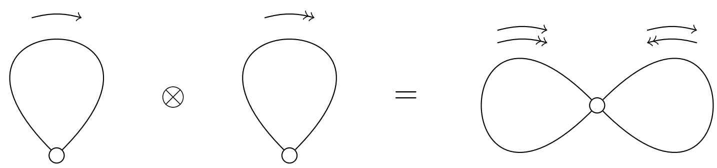

Figure 5.2: tensor product of two loops

times a loop creates two loops around the same vertex. Since we will need such direct product in Chapter 7, we take a different name and symbol for our generalized product.

There are some reasons to define tensor product in this manner. First, note the number of vertices in $L(X, v)$ is exactly the valency of v in X, where a loop at v contributes 2 to the number. Assuming $L(X, v) \bullet L(X', v') \cong L(X * X', (v, v'))$, this implies

$$d_X(v) \cdot d_{X'}(v') = d_{X*X'}((v, v')),$$

which is true for the tensor product, but not for the direct product admitting loops. Secondly, when we glue a face along a loop, the orientation of gluing matters, and the tensor product can keep track of such orientations. In Definition 5.1, when e_α or e_β is a loop, we shall think of it as an edge joining two different ends of the loop, say $+$ and $-$, and label two ends of $e^\delta_{\alpha,\beta}$ by $+$ and $-$ accordingly. We can then lift any given orientation of a loop in a factor to edges generated by this loop in the product, as illustrated in Figures 5.1 and 5.2. This also allows us to define projections unambiguously. Note that we do not assume graphs to be directed. We just distinguish two ends of each loop.

Definition 5.2. Assume the notation of Definition 5.1. The **projection** from $\Gamma \otimes \Gamma'$ to Γ, denoted by π_Γ, is a continuous function such that π_Γ maps $(v, v') \in V(\Gamma \otimes \Gamma')$ to $v \in V(\Gamma)$, and $e^\delta_{\alpha,\beta} \in E(\Gamma \otimes \Gamma')$ to $e_\alpha \in E(\Gamma)$ isometrically between endpoints. The projection $\pi_{\Gamma'}$ from $\Gamma \otimes \Gamma'$ to Γ' is likewise defined.

The projections defined above are graph homomorphisms in the following sense.

Definition 5.3. Let Γ and Γ' be two arbitrary graphs. A continuous function φ from Γ to Γ' is a **homomorphism** if φ maps each vertex of Γ to a vertex of Γ', and each open edge of Γ isometrically onto an open edge of Γ'.

Remark. In the above definition, the continuity of φ is essentially saying that a homomorphism maps incident vertices and edges to incident vertices and edges. Meanwhile, the isometric condition helps to choose a representative from all homotopic maps.

$$\begin{tikzcd}
 & \Gamma_0 \arrow[dl, "\varphi"'] \arrow[d, dashed, "\exists!\psi"'] \arrow[dr, "\varphi'"] & \\
\Gamma & \Gamma \otimes \Gamma' \arrow[l, "\pi_\Gamma"'] \arrow[r, "\pi_{\Gamma'}"] & \Gamma'
\end{tikzcd}$$

Figure 5.3: universal property of graph tensor product

Note that the composition of two graph homomorphisms is again a graph homomorphism. Together with the trivial automorphisms, the class of graphs forms a category. The following proposition shows that the tensor product defined above is actually the categorical product of this category.

Proposition 5.4. Let Γ and Γ' be two arbitrary graphs. Suppose that Γ_0 is a graph with two homomorphisms $\varphi : \Gamma_0 \to \Gamma$ and $\varphi' : \Gamma_0 \to \Gamma'$. Then there exists a unique homomorphism $\psi : \Gamma_0 \to \Gamma \otimes \Gamma'$ such that $\varphi = \pi_\Gamma \circ \psi$ and $\varphi' = \pi_{\Gamma'} \circ \psi$. In other words, there exists a unique ψ such that the diagram in Figure 5.3 commutes.

Proof. Assume that there exists a continuous function $\psi : \Gamma_0 \to \Gamma \otimes \Gamma'$ such that $\varphi = \pi_\Gamma \circ \psi$ and $\varphi' = \pi_{\Gamma'} \circ \psi$. Then $\forall v \in V(\Gamma_0)$, we have $\varphi(v) = \pi_\Gamma \circ \psi(v)$ and $\varphi'(v) = \pi_{\Gamma'} \circ \psi(v)$. By Definition 5.2, we know that $\psi(v) = (\varphi(v), \varphi'(v))$.

Suppose that e is an open edge joining v and u in Γ_0, and we denote $\varphi(e)$ and $\varphi'(e)$ by e_α and e_β respectively. By the continuity of ψ, $\psi(e)$ is an open path connecting $(\varphi(v), \varphi'(v))$ and $(\varphi(u), \varphi'(u))$. Notice that $e_\alpha = \varphi(e) = \pi_\Gamma \circ \psi(e)$ and $e_\beta = \varphi'(e) = \pi_{\Gamma'} \circ \psi(e)$. By Definition 5.2, we know $\psi(e)$ is either $e^0_{\alpha,\beta}$ or $e^1_{\alpha,\beta}$, determined by endpoints $(\varphi(v), \varphi'(v))$ and $(\varphi(u), \varphi'(u))$. In case e_α or e_β is a loop, by keeping track of ends of the loop, $\psi(e)$ is also uniquely determined. Moreover, the local isometry over open edges of φ and π_Γ forces ψ to map e isometrically to $\psi(e)$. Note that we have explicitly constructed a continuous ψ satisfying our initial assumption. We have also shown that ψ is uniquely determined, and actually a homomorphism, which finishes the proof. $\square$

For any two graphs Γ and Γ', we denote the set of all homomorphisms from Γ to Γ' by $\mathrm{Hom}(\Gamma, \Gamma')$. We have the following corollary about the number of homomorphisms.

Corollary 5.5. For any graphs Γ, Γ_1, Γ_2, we have

$$|\,\mathrm{Hom}(\Gamma, \Gamma_1 \otimes \Gamma_2)\,| = |\,\mathrm{Hom}(\Gamma, \Gamma_1))\,| \cdot |\,\mathrm{Hom}(\Gamma, \Gamma_2)\,|.$$

Proof. An immediate consequence of Proposition 5.4. $\square$

Note that for any graph Γ, there is a homomorphism from Γ to a loop. Since we distinguish the orientations when we map an edge to a loop, there are actually 2^n such homomorphisms, where n is the number of edges of Γ. In particular, a loop is not the terminal object in the category of arbitrary graphs.

Corollary 5.6. Let Γ and Γ' be two graphs, P be a path in Γ of length n from v to u, and P' be a path in Γ' of length n from v' to u'. Then in $\Gamma \otimes \Gamma'$, there exists a unique path, denoted by $(P, P')_{\otimes}$, from (v, v') to (u, u') such that $\pi_{\Gamma}((P, P')_{\otimes}) = P$ and $\pi_{\Gamma'}((P, P')_{\otimes}) = P'$.

Proof. Let I be a graph which is a path of length n. We can give I a specific orientation from one end to the other. Then there is a natural homomorphism φ from I to P, as well as one φ' from I to P'. By Proposition 5.4, there exists a unique homomorphism $\psi : I \to \Gamma \otimes \Gamma'$ such that $\varphi = \pi_{\Gamma} \circ \psi$ and $\varphi' = \pi_{\Gamma'} \circ \psi$. Hence we have $P = \varphi(I) = \pi_{\Gamma} \circ \psi(I)$ and $P' = \varphi'(I) = \pi_{\Gamma'} \circ \psi(I)$. Note that $\psi(I)$ satisfies the conditions of $(P, P')_{\otimes}$, and the uniqueness of $(P, P')_{\otimes}$ follows the uniqueness of ψ. $\square$

Remark. For simple graphs, this result is straightforward from the definition of tensor product. This corollary clarifies the case when P or P' contains a loop, where the orientation going through the loop will determine the edge to choose in $(P, P')_{\otimes}$.

To define our first complex product more concisely, we would like to extend the notation $(\ ,\ )_{\otimes}$ above. Let Γ_1 and Γ_2 be two graphs, C_1 be a cycle of length n in Γ_1, and C_2 be a cycle of length m in Γ_2. Both C_1 and C_2 are assigned initial vertices and orientations. Specifically, C_2 is $(v_0, e_0, v_1, e_1, \ldots, e_{m-1}, v_m = v_0)$, where $v_i \in V(\Gamma_2)$ and $e_j \in E(\Gamma_2)$. Then for $i \in \{0, 1, \ldots, m-1\}$ we define

$$(C_1, C_2)_{\otimes}^{i^{\delta}} := (\frac{[n,m]}{n} C_1, \frac{[n,m]}{m} C_2^{i^{\delta}})_{\otimes},$$

a cycle of length $[n, m]$ in $\Gamma_1 \otimes \Gamma_2$, where $[n, m]$ is the least common multiple of n and m, kC_j is the cycle repeating C_j k times, and $C_2^{i^0}$ is the same cycle as C_2, but starting at v_i, while $C_2^{i^1}$ is the reversed cycle of C_2 starting at v_i.

Definition 5.7. Let X and Y be two polygonal cell complexes with face sets $F(X) = \{f_{\alpha} \mid \alpha \in A\}$ and $F(Y) = \{f_{\beta} \mid \alpha \in B\}$. We denote the boundary length of f_{α} and f_{β} by n_{α} and n_{β} respectively, and let (n_{α}, n_{β}) denote the greatest common divisor of n_{α} and

n_β. The **tensor product** of X and Y, denoted by $X \otimes Y$, is a polygonal cell complex with 1-skeleton $X^1 \otimes Y^1$, the tensor product of the 1-skeletons of X and Y, and face set

$$F(X \otimes Y) = \{f_{\alpha,\beta}^{i^\delta} \mid \alpha \in A, \beta \in B, i \in \{0, 1, \ldots, (n_\alpha, n_\beta) - 1\}, \delta \in \{0, 1\}\},$$

where $f_{\alpha,\beta}^{i^\delta}$ is a face attached along $(C_\alpha, C_\beta)_\otimes^{i^\delta}$, while C_α is the cycle along which f_α is attached in X, and C_β is the cycle along which f_β is attached in Y.

Remark. We will use the jargon that $f_{\alpha,\beta}^{i^\delta}$ is generated by f_α and f_β, especially when faces are not clearly indexed. In the above definition, note that $(C_\alpha, C_\beta)_\otimes^{i^\delta}$ and $(C_\alpha, C_\beta)_\otimes^{i+(n_\alpha,n_\beta)^\delta}$ are identical cycles with different starting vertices. To let a pair of corners of f_α and f_β contribute to exactly one face corner in $X \otimes Y$, we only choose $i \in \{0, 1, \ldots, (n_\alpha, n_\beta) - 1\}$. Here we discard repeated corner pairs, not faces in $X \otimes Y$ attached along the same cycle. For example, let X and Y be 15-gons wrapped around a cycle of length 3 and 5 respectively. Note that the tensor product of a triangle and a pentagon is not the same as $X \otimes Y$. The former has only $2 \cdot (3, 5) = 2$ faces, while $X \otimes Y$ has $2 \cdot (15, 15) = 30$ faces in two groups, each of which has 15 faces with cyclically identical attaching maps.

In the example of a triangle tensor a pentagon, the only two faces meet at every vertex in the product. In general, when $n_\alpha \neq n_\beta$, two faces $f_{\alpha,\beta}^{i^0}$ and $f_{\alpha,\beta}^{i^1}$ meet at more than one vertex. Therefore the tensor product of two polygonal complexes is not necessarily polygonal. How about the case when $n_\alpha = n_\beta$? For n_α even, note that $f_{\alpha,\beta}^{i^0}$ and $f_{\alpha,\beta}^{i^1}$ have two vertices $(0, 0)$ and $(\frac{n_\alpha}{2}, \frac{n_\alpha}{2})$ in common, and the tensor product is not polygonal. For odd cases, we have the following result.

Proposition 5.8. Suppose that X and Y are polygonal complexes with all faces of the same odd length n. Then the tensor product $X \otimes Y$ is a polygonal complex.

Proof. Since X and Y are polygonal complexes, we know that X^1 and Y^1 are simple graphs, and hence the 1-skeleton of $X \otimes Y$, namely $X^1 \otimes Y^1$, is a simple graph as well. Consider the boundary of an arbitrary face $f_{\alpha,\beta}^{i^\delta}$ in $X \otimes Y$, namely $(C_\alpha, C_\beta)_\otimes^{i^\delta}$. Note that C_α and C_β are both simple closed cycles of the same length n, as they are boundaries of faces of polygonal complexes. Therefore $(C_\alpha, C_\beta)_\otimes^{i^\delta}$ is a simple closed cycle of length n. In brief, every face of $X \otimes Y$ is attached along a simple closed cycle.

Now all we have to show is that the intersection of two faces in $X \otimes Y$ is either empty, a vertex, or an edge in $X \otimes Y$. Suppose that there exist two faces $f_{\alpha,\beta}^{i^\delta}$ and $f_{\alpha',\beta'}^{j^{\delta'}}$ in $X \otimes Y$ such that the intersection of $f_{\alpha,\beta}^{i^\delta}$ and $f_{\alpha',\beta'}^{j^{\delta'}}$ is neither empty, a vertex, nor an edge. For the case of $n = 3$, it is not hard to see that $f^{i^\delta}_{\alpha,\beta}$ and $f^{j^{\delta'}}_{\alpha',\beta'}$ share the same boundary, and in fact are the same face by the polygonality of X and Y. For the case of odd $n > 3$, note that $f^{i^\delta}_{\alpha,\beta}$ and $f^{j^{\delta'}}_{\alpha',\beta'}$ share two vertices which are not consecutive on the boundary of faces. By the polygonality of X and Y, this implies that $f_\alpha = f_{\alpha'}$ and $f_\beta = f_{\beta'}$. Consider the boundaries of $f^{i^\delta}_{\alpha,\beta}$ and $f^{j^{\delta'}}_{\alpha,\beta}$, namely $(C_\alpha, C_\beta)^{i^\delta}_\otimes$ and $(C_\alpha, C_\beta)^{j^{\delta'}}_\otimes$. When $\delta = \delta'$ and $i \neq j$, $(C_\alpha, C_\beta)^{i^\delta}_\otimes$ and $(C_\alpha, C_\beta)^{j^{\delta'}}_\otimes$ have no vertex in common. When $\delta \neq \delta'$, notice that a common vertex of $(C_\alpha, C_\beta)^{i^\delta}_\otimes$ and $(C_\alpha, C_\beta)^{j^{\delta'}}_\otimes$ corresponds to an integer m such that

$$j + m \equiv i - m \mod n \;\Leftrightarrow\; 2m = i - j \mod n,$$

which has a unique solution when n is odd. In other words, when $\delta \neq \delta'$, $(C_\alpha, C_\beta)^{i^\delta}_\otimes$ and $(C_\alpha, C_\beta)^{j^{\delta'}}_\otimes$ intersect at exactly one vertex. Since $f^{i^\delta}_{\alpha,\beta}$ and $f^{j^{\delta'}}_{\alpha,\beta}$ share two vertices, we can conclude that $\delta = \delta'$ and $i = j$. This finishes the proof. □

The complex tensor product does not preserve simple connectedness either.

Proposition 5.9. Let X and Y be an n-gon and m-gon respectively, where n and m are two positive integers. Then $X \otimes Y$ is simply-connected if and only if $n = m = 1$.

Proof. When $n = m = 1$, the 1-skeleton of $X \otimes Y$ is a vertex with two loops, as illustrated in Figure 5.2, and $X \otimes Y$ has two faces attached along these two loops respectively. In this case, $X \otimes Y$ is actually contractible, and of course simply-connected.

Now suppose that n and m are not both equal to 1. Without loss of generality, we can assume $n \geq 2$. Note that X has n vertices, n edges, and 1 face, whereas Y has m vertices, m edges, and 1 face. By Definition 5.7, the complex $X \otimes Y$ has nm vertices, $2nm$ edges, and $2(n, m)$ faces. Therefore $X \otimes Y$ has Euler characteristic

$$\chi(X \otimes Y) = nm - 2nm + 2(n, m) = -nm + 2(n, m) \leq -2m + 2m = 0.$$

By Proposition 2.2, we know $X \otimes Y$ is not simply-connected. □

Remark. Let X and Y be two arbitrary complexes, and C be a cycle along the 1-skeleton of $X \otimes Y$. This proposition shows that the contractibility of $\pi_X(C)$ and $\pi_Y(C)$ does not guarantee the contractibility of C. Conversely, when C is contractible in $X \otimes Y$, can we conclude that $\pi_X(C)$ and $\pi_Y(C)$ are contractible? The answer is positive. We can find a series $\{C_j\}$ of homotopic cycles of C such that $C_0 = C$, C_n is a vertex, and each C_j morphs through a single face $f^{i^\delta}_{\alpha,\beta}$ to obtain C_{j+1}. Note that $\pi_X(C_j)$ can morph through a single face f_α to obtain $\pi_X(C_{j+1})$, even when the length of f_α properly divides the length of $f^{i^\delta}_{\alpha,\beta}$. Therefore $\pi_X(C) = \pi_X(C_0)$ is homotopic to $\pi_X(C_n)$, which is a vertex.

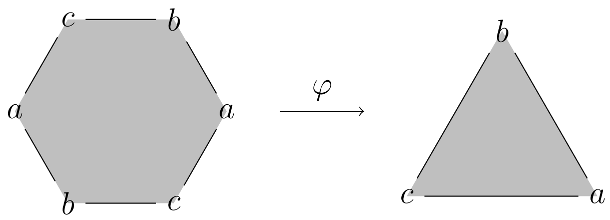

Figure 5.4: a complex homomorphism from a hexagon to a triangle

In the above remark, we actually abuse the notation π_X, as we have not yet defined projection maps for complex tensor products. To define such projection maps, first we introduce some terminology. Let X and Y be an n-gon and m-gon with centre O_X and O_Y respectively. A function $\rho : X \to Y$ is **radial** if ρ sends O_X to O_Y, ∂X to ∂Y, and for every point $P \in \partial X$, every real number $t \in [0, 1]$, we have

$$\rho\big(t \cdot O_X + (1-t)P\big) = t \cdot O_Y + (1-t)\rho(P).$$

Definition 5.10. Assume the notation of Definition 5.7. The **projection** from $X \otimes Y$ to X, denoted by π_X, is a continuous function such that π_X restricted to $X^1 \otimes Y^1$ is exactly π_{X^1}, the projection of the graph tensor product, and π_X maps $f^{i^\delta}_{\alpha,\beta} \in F(X \otimes Y)$ radially to $f_\alpha \in F(X)$. The projection π_Y from $X \otimes Y$ to Y is likewise defined.

The projection maps defined above are complex homomorphisms in the following sense.

Definition 5.11. Let X and Y be two polygonal cell complexes. A continuous function φ from X to Y is a **homomorphism** if φ restricted to X^1 is a graph homomorphism to Y^1, and φ maps each face of X radially to a face of Y and each open face corner (ignoring the boundary) of X homeomorphically to an open face corner of Y.

Remark. In the above definition, the continuity of φ is essentially saying that a complex homomorphism maps incident cells to incident cells. Similar to the isometric condition in graph homomorphism, the radial condition is imposed to rule out homotopic complex homomorphisms. Most important of all, the homeomorphic corner condition forces a face of X to wrap around a face f of Y along the direction of the attaching map of f, possibly more that once. In particular, a face of length n can only be mapped to a face of length dividing n. Figure 5.4 illustrates such phenomenon, where corners are mapped to a corner with the same label. The projection π_X of complex tensor product mapping $f^{i^\delta}_{\alpha,\beta}$ to f_α is also a typical example.

$$\begin{tikzcd}
& Z \arrow[dl, "\varphi_X"'] \arrow[d, dashed, "\exists!\psi"'] \arrow[dr, "\varphi_Y"] & \\
X & X \otimes Y \arrow[l, "\pi_X"'] \arrow[r, "\pi_Y"] & Y
\end{tikzcd}$$

Figure 5.5: universal property of complex tensor product

Note that the composition of two complex homomorphisms is again a complex homomorphism. Together with the trivial automorphisms, the class of polygonal cell complexes forms a category. The following proposition shows that the complex tensor product defined above is actually the categorical product of this category.

Proposition 5.12. Let X and Y be two polygonal cell complexes. Suppose that Z is a complex with two homomorphisms $\varphi_X : Z \to X$ and $\varphi_Y : Z \to Y$. Then there exists a unique homomorphism $\psi : Z \to X \otimes Y$ such that $\varphi_X = \pi_X \circ \psi$ and $\varphi_Y = \pi_Y \circ \psi$. In other words, there exists a unique ψ such that the diagram in Figure 5.5 commutes.

Proof. Assume that there exists a continuous function $\psi : Z \to X \otimes Y$ such that $\varphi_X = \pi_X \circ \psi$ and $\varphi_Y = \pi_Y \circ \psi$. Note that φ_X, φ_Y, π_X, and π_Y restricted to the 1-skeletons of their domains are all graph homomorphisms. By Proposition 5.4, the restriction of ψ to Z^1 is a uniquely determined graph homomorphism to $X^1 \otimes Y^1$.

Suppose that f is a face in Z, $\varphi_X(f)$ wraps around a face f_α in X, and $\varphi_Y(f)$ wraps around a face f_β in X. Then $\varphi_X(f) = \pi_X \circ \psi(f)$ wraps around f_α, and $\varphi_Y(f) = \pi_Y \circ \psi(f)$ wraps around f_β. By Definition 5.10, $\psi(f)$ must wrap around $f^{i^\delta}_{\alpha,\beta}$ for some i and δ. Let c be a corner of f. Then we must have $\varphi_X(c) = \pi_X \circ \psi(c)$ and $\varphi_Y(c) = \pi_Y \circ \psi(c)$. By the remark after Definition 5.7, this pair of corners $(\varphi_X(c), \varphi_Y(c))$, orientation included, appears in exactly one $f^{i^\delta}_{\alpha,\beta}$. Therefore i and δ are uniquely determined, and $\psi(f)$ wraps around this $f^{i^\delta}_{\alpha,\beta}$. Moreover, the radiality of φ_X and π_X forces ψ to map f radially to $f^{i^\delta}_{\alpha,\beta}$. Note that we have explicitly constructed a continuous ψ satisfying our initial assumption. We have also shown that ψ is uniquely determined, and actually a complex homomorphism, which finishes the proof. $\square$

Remark. For any two complexes X and Y, we denote the set of all complex homomorphisms from X to Y by $\mathrm{Hom}(X, Y)$. Similarly to Corollary 5.5, we have

$$|\,\mathrm{Hom}(Z, X \otimes Y)\,| = |\,\mathrm{Hom}(Z, X)\,| \cdot |\,\mathrm{Hom}(Z, Y)\,|.$$

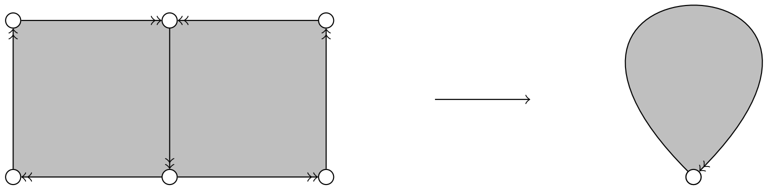

Figure 5.6: a homomorphism to a 1-gon

$$\begin{tikzcd}
& Y \arrow[dr, "\rho"] & & & & & L(Y,u) \arrow[dr, "L(\rho)"] & \\
X \arrow[ur, "\varphi"] \arrow[rr, "\rho\circ\varphi"'] & & Z & {} \arrow[r, "L"] & {} & L(X,v) \arrow[ur, "L(\varphi)"] \arrow[rr, "L(\rho\circ\varphi)"'] & & L(Z,w)
\end{tikzcd}$$

Figure 5.7: functoriality of L

As we mentioned earlier, for any graph Γ, there is a homomorphism from Γ to a loop. It is reasonable to ask the following question: for any complex X, is there always a homomorphism from X to a 1-gon? The answer is negative. Take Figure 5.6 as an example. Once the image of the leftmost edge is determined, it determines the image of all other edges. If we identify the leftmost and the rightmost edges with a twist, i.e. making it a Mobius strip, then there is no way to have a homomorphism. Note that this question is not related to orientability. If the complex is a strip with 3 squares, then the Mobius case has a homomorphism, while the orientable case does not.

Proposition 5.13. Let X and Y be two polygonal cell complexes, and $\varphi : X \to Y$ be a complex homomorphism mapping a vertex $v \in V(X)$ to $u \in V(Y)$. Then φ induces a graph homomorphism $L(\varphi)$ from $L(X, v)$ to $L(Y, u)$. Moreover, let Z be another complex and $\rho : Y \to Z$ be a complex homomorphism mapping u to $w \in V(Z)$. Then we have $L(\rho \circ \varphi) = L(\rho) \circ L(\varphi)$, as illustrated in Figure 5.7.

Proof. By Definition 1.6, $L(X, v)$ has vertices corresponding to edge ends around v in X, and edges corresponding to face corners at v in X. Since φ restricted to X^1 is a graph homomorphism, φ maps an edge end around v in X to an edge end around u in Y. In addition, by the homeomorphic condition in Definition 5.11, φ maps a face corner at v joining two edge ends around v homeomorphically to a face corner at u joining two

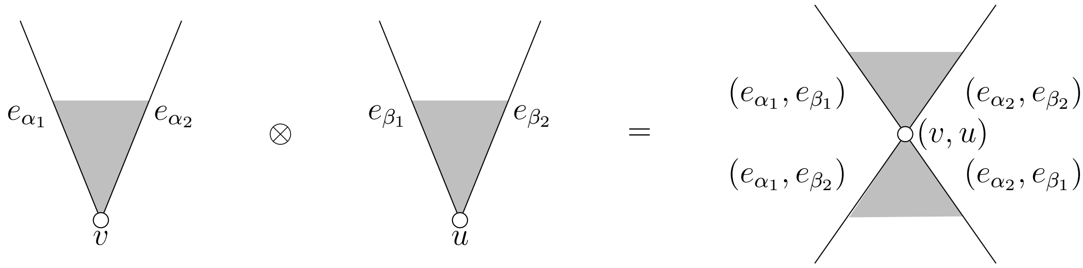

Figure 5.8: well linked tensor product

edge ends around u. Therefore φ induces a graph homomorphism $L(\varphi)$ from $L(X, v)$ to $L(Y, u)$. Once these induced graph homomorphisms between link graphs are defined, the equality $L(\rho \circ \varphi) = L(\rho) \circ L(\varphi)$ follows immediately. □

Remark. To each polygonal cell complex, we can assign a distinguished vertex to be the basepoint. Together with basepoint-preserving homomorphisms, the class of pointed polygonal cell complexes also forms a category. The above proposition is essentially saying that L is a functor from this category to the category of graphs.

Now we move back to the main purpose of this chapter: to develop a complex product interacting nicely with some product of link graphs. From the above discussion, we know that the complex tensor product arises naturally in the category of polygonal cell complexes. Does this natural categorical product fulfill the main job? Yes, it does.

Theorem 5.14. Suppose that X and Y are two polygonal cell complexes, and v and u are two vertices in X and Y respectively. Then we have

$$L(X, v) \otimes L(Y, u) \cong L(X \otimes Y, (v, u)).$$

Proof. We can identify edge ends incident to a vertex as paths of length 1 leaving the vertex, since a loop contributes to two edge ends as well as two such paths, which we call 1-paths for short. By Corollary 5.6, there is a bijection between 1-paths leaving (v, u) in $X \otimes Y$ and pairs of 1-path leaving v in X and 1-path leaving u in Y. Therefore we can index 1-paths leaving (v, u) in $X \otimes Y$ by such 1-path pairs in X and Y.

Suppose that $f_\alpha \in F(X)$ has a corner c_α at $(e_{\alpha_1}, v, e_{\alpha_2})$, and $f_\beta \in F(Y)$ has a corner c_β at $(e_{\beta_1}, u, e_{\beta_2})$, as illustrated in Figure 5.8. These e_*'s should be understood as 1-paths. By the remark after Definition 5.7, the pairing of these two corners appears exactly once in $f^{i^0}_{\alpha,\beta}$ and $f^{j^1}_{\alpha,\beta}$ respectively, forming corners $((e_{\alpha_1}, e_{\beta_1}), (v, u), (e_{\alpha_2}, e_{\beta_2}))$ and

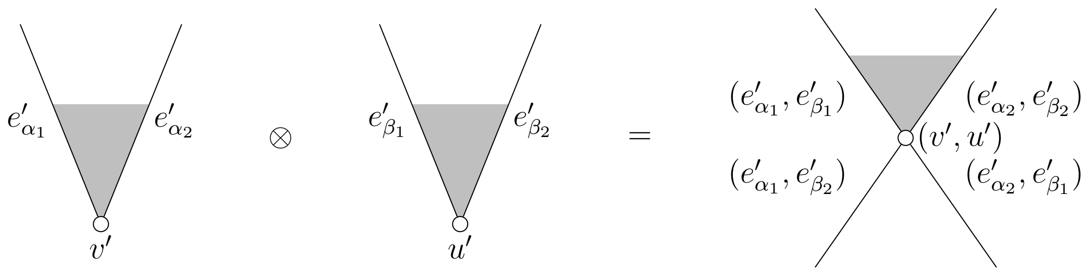

Figure 5.9: automorphic image of Figure 5.8

$((e_{\alpha_1}, e_{\beta_2}), (v, u), (e_{\alpha_2}, e_{\beta_1}))$ in $X \otimes Y$. Note that by taking projection maps, we know that any face corner at (v, u) comes from some pairing of corners at v and u.

Now we translate the above statements in terms of corresponding link graphs. First of all, we have $V\big(L(X, v)\big) \times V\big(L(Y, u)\big) \cong V\big(L(X \otimes Y, (v, u))\big)$. Secondly, the corner c_α is an edge joining vertices e_{α_1} and e_{α_2} in $L(X, v)$, and c_β is an edge joining vertices e_{β_1} and e_{β_2} in $L(Y, u)$. Notice that the edge pair (c_α, c_β) contributes to one edge joining $(e_{\alpha_1}, e_{\beta_1})$ and $(e_{\alpha_2}, e_{\beta_2})$, and one edge joining $(e_{\alpha_1}, e_{\beta_2})$ and $(e_{\alpha_2}, e_{\beta_1})$ in $L(X \otimes Y, (v, u))$. Meanwhile, taking all possible pairings of edges exhausts all edges in $L(X \otimes Y, (v, u))$. By Definition 5.1, this is exactly saying that $L(X, v) \otimes L(Y, u) \cong L(X \otimes Y, (v, u))$. $\square$

Remark. In the terminology of category theory, this theorem is essentially saying that the functor L from the category of pointed complexes to the category of graphs preserves categorical products, which is not always true for an arbitrary functor.

As indicated in Propositions 5.8 and 5.9, the complex tensor product does not necessarily preserve polygonality and simple connectedness. Fortunately, complex tensor product does preserve the most important property for our purpose.

Theorem 5.15. Let X and Y be any platonic polygonal cell complexes. Then the complex tensor product $X \otimes Y$ is also a platonic complex.

Proof. In case X or Y has no faces, then $X \otimes Y$ is simply a graph, and the platonicity follows easily from the definition of graph tensor product. Hereafter we assume that both X and Y have at least one face. Since X and Y are platonic, X and Y have the property that each vertex is incident to an edge, and each edge is incident to a face. By Definition 5.7, it is easy to see that $X \otimes Y$ has the same property as well. Hence to show $X \otimes Y$ is platonic, it suffices to show that $X \otimes Y$ is flag-transitive.

Let $((e_{\alpha_1}, e_{\beta_1}), (v, u), (e_{\alpha_2}, e_{\beta_2}))$ be a face corner in $X \otimes Y$, which projects to a corner $(e_{\alpha_1}, v, e_{\alpha_2})$ in X and a corner $(e_{\beta_1}, u, e_{\beta_2})$ in Y, as illustrated in Figure 5.8. Let

$((e'_{\alpha_1}, e'_{\beta_1}), (v', u'), (e'_{\alpha_2}, e'_{\beta_2}))$ be another face corner in $X \otimes Y$, which projects to a corner $(e'_{\alpha_1}, v', e'_{\alpha_2})$ in X and a corner $(e'_{\beta_1}, u', e'_{\beta_2})$ in Y, as illustrated in Figure 5.9. Since X and Y are platonic, there exist $\rho \in \mathrm{Aut}(X)$ mapping $(e_{\alpha_1}, v, e_{\alpha_2})$ to $(e'_{\alpha_1}, v', e'_{\alpha_2})$ and $\sigma \in \mathrm{Aut}(Y)$ mapping $(e_{\beta_1}, u, e_{\beta_2})$ to $(e'_{\beta_1}, u', e'_{\beta_2})$. Comparing Figures 5.8 and 5.9, note that (ρ, σ) gives an automorphism of $X \otimes Y$ mapping $((e_{\alpha_1}, e_{\beta_1}), (v, u), (e_{\alpha_2}, e_{\beta_2}))$ to $((e'_{\alpha_1}, e'_{\beta_1}), (v', u'), (e'_{\alpha_2}, e'_{\beta_2}))$. The above discussion shows that $\mathrm{Aut}(X \otimes Y)$ acts transitively on face corners with orientations, and therefore transitively on half-corners. In other words, $\mathrm{Aut}(X \otimes Y)$ acts transitively on flags, and $X \otimes Y$ is a platonic complex. □

Remark. In Figure 5.8, flipping both corners in X and Y will flip both corners in $X \otimes Y$, whereas flipping only one corner in either X or Y will swap two corners in $X \otimes Y$.

Now we can easily construct platonic complexes with $B_{n,n}$ links.

Corollary 5.16. Let X be a platonic polygonal cell complex with K_n links, and Y be a polygon of arbitrary length. Then $X \otimes Y$ is a platonic complex with $B_{n,n}$ links.

Proof. Note that the polygon Y is a platonic complex with K_2 links. By Theorem 5.15, the tensor product $X \otimes Y$ is a platonic complex. By Theorem 5.14, each vertex in $X \otimes Y$ has link graph $K_n \otimes K_2 = B_{n,n}$. This completes the proof. □

In the proof of Theorem 5.15, the key fact we used is the following relation:

$$\mathrm{Aut}(X) \times \mathrm{Aut}(Y) \leq \mathrm{Aut}(X \otimes Y).$$

Is it possible that these two groups are actually isomorphic? We will answer this question in Chapter 7, where the symmetry of tensor products is studied in greater detail.

Chapter 6

Zigzag Product

In the previous chapter, we constructed platonic complexes with $B_{n,n}$ links using the complex tensor product, based on the fact that $B_{n,n} = K_n \otimes K_2$. In this chapter, we will develop a different product of complexes to construct platonic complexes with Q_n links, which is related to another product of graphs.

Let Γ and Γ' be two arbitrary simple graphs. The **Cartesian product** of Γ and Γ', denoted by $\Gamma \square \Gamma'$, is a simple graph with vertex set $V(\Gamma) \times V(\Gamma')$. Two vertices (v, v') and (u, u') are adjacent in $\Gamma \square \Gamma'$ if and only if one of the following is true:

(1) $v = u$ and v' is adjacent to u' in Γ'.

(2) $v' = u'$ and v is adjacent to u in Γ.

The definition can be easily understood through Figure 6.1, showing that the Cartesian product of two paths of length n and m results in an $n \times m$ Cartesian grid, which explains the name of the product. Note that the Cartesian product of two edges is a 1×1 grid, which explains the symbol "$\square$" of the product.

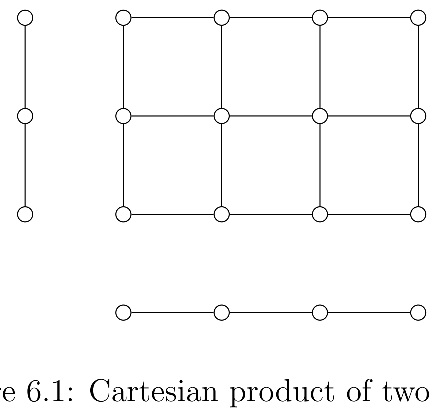

Figure 6.1: Cartesian product of two paths

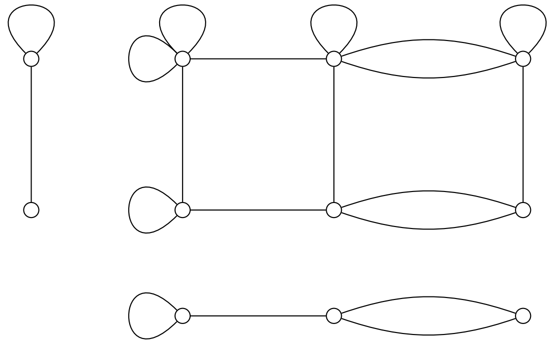

Figure 6.2: Cartesian product of non-simple graphs

By induction, it is straightforward to show that the Cartesian product of graphs is associative, the product of $\Gamma_1, \Gamma_2, \ldots, \Gamma_n$, denoted by $\square_{i=1}^n \Gamma_i$, has vertex set $\times_{i=1}^n V(\Gamma_i)$, and two vertices $(v_1, \ldots, v_n)$ and $(u_1, \ldots, u_n)$ are adjacent in $\square_{i=1}^n \Gamma_i$ if and only if $\exists j \in \{1, \ldots, n\}$ such that v_j and u_j are adjacent in Γ_j, and $v_i = u_j$ for any other $i \neq j$. In particular, when each $\Gamma_i = K_2$, the resulting graph $\square_{i=1}^n K_2$ is exactly the hypercube graph Q_n. This motivates us to develop a complex product $\diamond$ such that

$$L(X, v) \square L(Y, u) \cong L(X \diamond Y, (v, u)).$$

We would like to have this property for arbitrary complexes and links. Therefore we extend the definition of Cartesian product to non-simple graphs first.

Definition 6.1. Let Γ_1 and Γ_2 be two arbitrary graphs with $V(\Gamma_j) = \{v_{\alpha_j} \mid \alpha_j \in V_j\}$ and $E(\Gamma_j) = \{e_{\overline{\alpha_j}} \mid \overline{\alpha_j} \in E_j\}$. Then the **Cartesian product** of Γ_1 and Γ_2, denoted by $\Gamma_1 \square \Gamma_2$, is a graph with vertex set $V(\Gamma_1 \square \Gamma_2) = V(\Gamma_1) \times V(\Gamma_2)$, and edge set

$$E(\Gamma_1 \square \Gamma_2) = \{e_{\overline{\alpha_1}, \alpha_2}, e_{\alpha_1, \overline{\alpha_2}} \mid \alpha_j \in V_j, \overline{\alpha_j} \in E_j\},$$

where $e_{\overline{\alpha_1}, \alpha_2}$ is an edge joining (v_1, v_{α_2}) and (u_1, v_{α_2}), given that $e_{\overline{\alpha_1}}$ joins v_1 and u_1 in Γ_1, while $e_{\alpha_1, \overline{\alpha_2}}$ is an edge joining (v_{α_1}, v_2) and (v_{α_1}, u_2), given that $e_{\overline{\alpha_2}}$ joins v_2 and u_2 in Γ_2.

Remark. The above lengthy definition is illustrated in Figure 6.2. The key point is that, for any vertex $v_{\alpha_2} \in \Gamma_2$, there is a corresponding copy of Γ_1 in $\Gamma_1 \square \Gamma_2$, similarly for $v_{\alpha_1} \in \Gamma_1$ and Γ_2. When restricted to simple graphs, the product defined above coincides with the ordinary Cartesian product. Note that $\forall v_{\alpha_1} \in V(\Gamma_1)$ and $\forall v_{\alpha_2} \in V(\Gamma_2)$, we have

$$d((v_{\alpha_1}, v_{\alpha_2})) = d(v_{\alpha_1}) + d(v_{\alpha_2}),$$

a valency relation which should be preserved under reasonable generalizations.

Like the ordinary Cartesian product, the generalized product is also associative.

Proposition 6.2. For $j \in \{1,2,3\}$, let Γ_j be a graph with $V(\Gamma_j) = \{v_{\alpha_j} \,|\, \alpha_j \in V_j\}$ and $E(\Gamma_j) = \{e_{\overline{\alpha_j}} \,|\, \overline{\alpha_j} \in E_j\}$. Suppose Γ is a graph with $V(\Gamma) = \{(v_{\alpha_1}, v_{\alpha_2}, v_{\alpha_3}) \,|\, \alpha_j \in V_j\}$ and

$$E(\Gamma) = \{e_{\overline{\alpha_1},\alpha_2,\alpha_3}, e_{\alpha_1,\overline{\alpha_2},\alpha_3}, e_{\alpha_1,\alpha_2,\overline{\alpha_3}} \mid \alpha_j \in V_j, \overline{\alpha_j} \in E_j\},$$

where $e_{\overline{\alpha_1},\alpha_2,\alpha_3}$ is an edge joining $(v_1, v_{\alpha_2}, v_{\alpha_3})$ and $(u_1, v_{\alpha_2}, v_{\alpha_3})$, given that $e_{\overline{\alpha_1}}$ joins v_1 and u_1 in Γ_1, while $e_{\alpha_1,\overline{\alpha_2},\alpha_3}$ and $e_{\alpha_1,\alpha_2,\overline{\alpha_3}}$ are similarly defined. Then we have

$$\Gamma \cong (\Gamma_1 \Box \Gamma_2) \Box \Gamma_3 \cong \Gamma_1 \Box (\Gamma_2 \Box \Gamma_3).$$

Proof. We will only show $\Gamma \cong (\Gamma_1 \Box \Gamma_2) \Box \Gamma_3$ here, and the same argument also applies to $\Gamma \cong \Gamma_1 \Box (\Gamma_2 \Box \Gamma_3)$. By Definition 6.1, there are two types of edges in $(\Gamma_1 \Box \Gamma_2) \Box \Gamma_3$:

(1) An edge corresponds to an edge, $e_{\overline{\alpha_1},\alpha_2}$ or $e_{\alpha_1,\overline{\alpha_2}}$, in $\Gamma_1 \Box \Gamma_2$ and a vertex v_{α_3} in Γ_3.

- $e_{\overline{\alpha_1},\alpha_2}$: Given that $\overline{\alpha_1}$ joins v_1 and u_1 in Γ_1, then $e_{\overline{\alpha_1},\alpha_2}$ joins (v_1, v_{α_2}) and (u_1, v_{α_2}) in $\Gamma_1 \Box \Gamma_2$. Hence this edge joins $((v_1, v_{\alpha_2}), v_{\alpha_3})$ and $((u_1, v_{\alpha_2}), v_{\alpha_3})$ in $(\Gamma_1 \Box \Gamma_2) \Box \Gamma_3$.
- $e_{\alpha_1,\overline{\alpha_2}}$: Given that $\overline{\alpha_2}$ joins v_2 and u_2 in Γ_2, then $e_{\alpha_1,\overline{\alpha_2}}$ joins (v_{α_1}, v_1) and (v_{α_1}, u_1) in $\Gamma_1 \Box \Gamma_2$. Hence this edge joins $((v_{\alpha_1}, v_1), v_{\alpha_3})$ and $((v_{\alpha_1}, u_1), v_{\alpha_3})$ in $(\Gamma_1 \Box \Gamma_2) \Box \Gamma_3$.

(2) An edge corresponds to a vertex $(v_{\alpha_1}, v_{\alpha_2})$ in $\Gamma_1 \Box \Gamma_2$ and an edge e_{α_3} in Γ_3. Given that e_{α_3} joins v_3 and u_3 in Γ_3, then this edge joins $\big((v_{\alpha_1}, v_{\alpha_2}), v_3\big)$ and $\big((v_{\alpha_1}, v_{\alpha_2}), u_3\big)$.

Now we identify the vertex $\big((v_{\alpha_1}, v_{\alpha_2}), v_{\alpha_3}\big)$ in $(\Gamma_1 \Box \Gamma_2) \Box \Gamma_3$ with the vertex $(v_{\alpha_1}, v_{\alpha_2}, v_{\alpha_3})$ in Γ. Then the above edges of different types correspond to $e_{\overline{\alpha_1},\alpha_2,\alpha_3}$, $e_{\alpha_1,\overline{\alpha_2},\alpha_3}$, and $e_{\alpha_1,\alpha_2,\overline{\alpha_3}}$ respectively, and this gives an isomorphism between $(\Gamma_1 \Box \Gamma_2) \Box \Gamma_3$ and Γ. $\Box$

We are about to define a complex product which interacts nicely with the graph Cartesian product. Note that for the graph Cartesian product, we have $V(\Gamma_1 \Box \Gamma_2) = V(\Gamma_1) \times V(\Gamma_2)$. According to the discussion in Chapter 5, the 1-skeleton of a desired complex product here must be the graph tensor product of 1-skeletons of factors. The remaining job is to attach faces, and we use the notation $(\ ,\ )_\otimes$ in Proposition 5.6 again. Let Γ_1 and Γ_2 be two graphs, C be a cycle of length n in Γ_1, and e be an edge joining v_0 and v_1 in Γ_2. Then we define

$$(C, e)^i_\otimes = (\frac{[n,2]}{n} C, \frac{[n,2]}{2} C^i_e)_\otimes,$$

where C^i_e is a cycle of length 2 going back and forth along e, while $i \in \{0,1\}$ indicates v_0 or v_1 to be the starting vertex. And $(e', C')^i_\otimes$ is similarly defined.

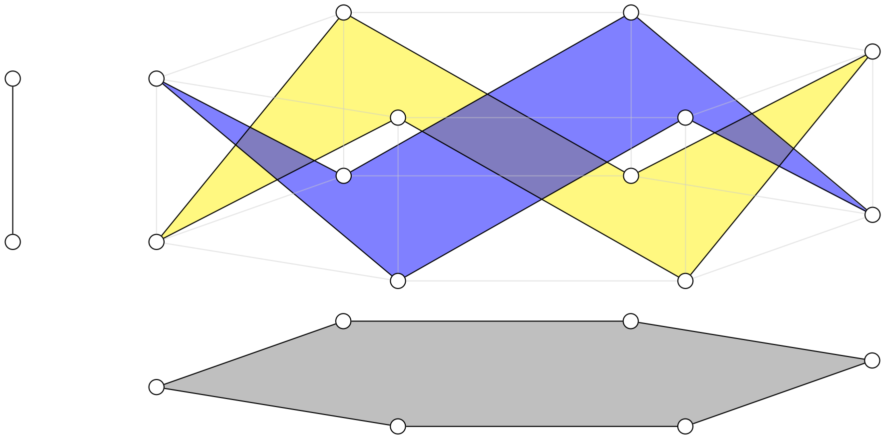

Figure 6.3: zigzag product of K_2 with a hexagon

Definition 6.3. For $j \in \{1,2\}$, let X_j be a polygonal cell complex with edge set $E(X_j) = \{e_{\alpha_j} \mid \alpha_j \in E_j\}$ and face set $F(X_j) = \{f_{\overline{\alpha_j}} \mid \overline{\alpha_j} \in F_j\}$. The **zigzag product** of X_1 and X_2, denoted by $X_1 \diamond X_2$, is a polygonal cell complex with 1-skeleton $X_1^1 \otimes X_2^1$, the tensor product of 1-skeletons of X_1 and X_2, and face set

$$F(X_1 \diamond X_2) = \left\{ f^{i_2}_{\overline{\alpha_1},\alpha_2}, f^{i_1}_{\alpha_1,\overline{\alpha_2}} \mid \alpha_j \in E_j, \overline{\alpha_j} \in F_j \right\},$$

where $f^{i_2}_{\overline{\alpha_1},\alpha_2}$ is a face attached along $(C_{\overline{\alpha_1}}, e_{\alpha_2})^{i_2}_{\otimes}$, while $C_{\overline{\alpha_1}}$ is the cycle of length $n_{\overline{\alpha_1}}$ along which $f_{\overline{\alpha_1}}$ is attached in X_1, and $i_2 \in \{0, (2, n_{\overline{\alpha_1}}) - 1\}$. The face $f^{i_1}_{\alpha_1,\overline{\alpha_2}}$ is similarly defined.

We shall explain $f^{i_1}_{\alpha_1,\overline{\alpha_2}}$ in words and figures to complement this complicated notation. An edge e_{α_1} in X_1 and a face $f_{\overline{\alpha_2}}$ in X_2 generate either one or two faces $f^{i_1}_{\alpha_1,\overline{\alpha_2}}$'s in $X_1 \diamond X_2$ depending on the parity of $n_{\overline{\alpha_2}}$, the boundary length of $f_{\overline{\alpha_2}}$. The face $f^{i_1}_{\alpha_1,\overline{\alpha_2}}$ is attached along a cycle in $X_1 \diamond X_2$ where the X_2 coordinate goes around the boundary of $f_{\overline{\alpha_2}}$, while the X_1 coordinate goes back and forth along e. When $n_{\overline{\alpha_2}}$ is even, it takes $n_{\overline{\alpha_2}}$ steps to return to the starting vertex, and two different starting vertices in X_1 create two different faces, as illustrated in Figure 6.3. When $n_{\overline{\alpha_2}}$ is odd, it takes $2n_{\overline{\alpha_2}}$ steps to return, and two different starting vertices in X_1 actually give the same cycle in $X_1 \diamond X_2$, so we merely choose one face $f^0_{\alpha_1,\overline{\alpha_2}}$ to attach, as illustrated in Figure 6.4.

We would like to clarify the case when e is a loop, illustrated in Figure 6.4 as well. Then going back and forth along e means going around the loop with alternating orientations. While a loop doubles edges in the product, it also induces orientations on these edges, and we attach a face generated by this loop along edges in alternating orientations. A good way to visualize this in Figure 6.4 is to vertically pinch together vertices of the 10-gon in K_2 times a pentagon, and this gives the 10-gon in a loop times a pentagon.

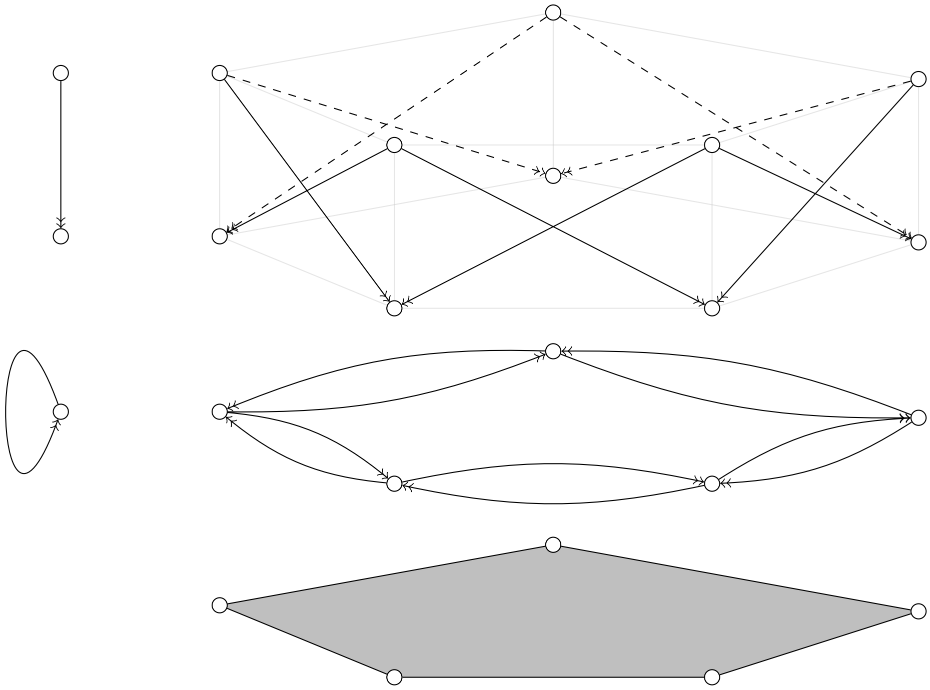

Figure 6.4: zigzag product of K_2 plus a loop with a pentagon

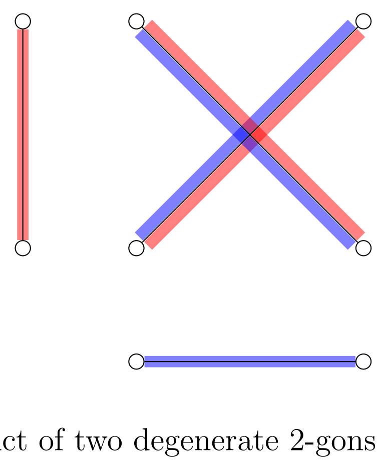

Figure 6.5: product of two degenerate 2-gons, tensor or zigzag?

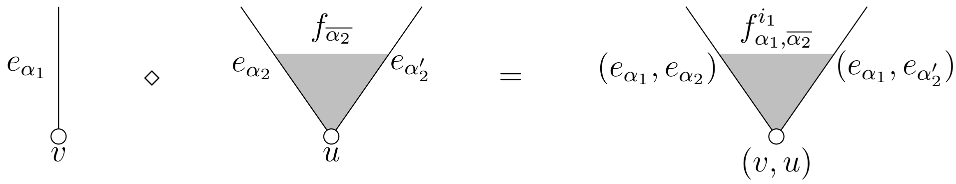

Figure 6.6: well linked zigzag product

The reader might have noticed the similarity between the zigzag product of an edge with a face and the tensor product of a degenerate 2-gon, a 2-gon attached on an edge, with a face. Note that in the tensor product of two polygons X and Y, we match corners of X and Y to generate faces in $X \otimes Y$, and in general flipping a corner of X gives a different face in $X \otimes Y$. In case X is a degenerate 2-gon, then flipping a corner of X results in an identical face. We keep such identical faces in the tensor product, but not in the zigzag product. In particular, the tensor product of a degenerate 2-gon with an n-gon has $2 \cdot (2, n)$ faces, while the zigzag product of an edge with an n-gon has $(2, n)$ faces.

The product of two degenerate 2-gons is an example of its own interest. In this example both tensor product and zigzag product generate two edges, each of which has two degenerate 2-gons attached, as illustrated in Figure 6.5, where thick strips indicate the presence of degenerate 2-gons. In the zigzag product case, the coloured strip in a factor is responsible for two strips of the same colour in the product. Note that the link of the product is a vertex with two loops attached, which is the graph tensor product of two loops, as well as the Cartesian product of two loops, namely the links of two factors.

Theorem 6.4. Suppose that X_1 and X_2 are two polygonal cell complexes, and v and u are two vertices in X_1 and X_2 respectively. Then we have

$$L(X_1, v) \square L(X_2, u) \cong L(X_1 \diamond X_2, (v, u)).$$

Proof. As in the proof of Theorem 5.14, we can identify edge ends incident to a vertex as 1-paths leaving the vertex, and by Corollary 5.6 we can index 1-paths leaving (v, u) in $X_1 \diamond X_2$ by pairs of 1-paths leaving v in X_1 and 1-path leaving u in X_2.

Assume the notation of Definition 6.3, and look at a face corner at (v, u) in $X_1 \diamond X_2$ generated by an edge $e_{\alpha_1} \in X_1$ and a face $f_{\overline{\alpha_2}} \in X_2$, as illustrated in Figure 6.6, where those e_*'s should be understood as 1-paths. By Definition 6.3, a corner of $f_{\overline{\alpha_2}}$ at u contributes to exactly one corner of $f^{i_1}_{\alpha_1,\overline{\alpha_2}}$ at (v, u), no matter whether the length of $f_{\overline{\alpha_2}}$ is even or odd. Note that a corner of $f_{\overline{\alpha_2}}$ joins 1-paths e_{α_2} and $e_{\alpha'_2}$ if and only if

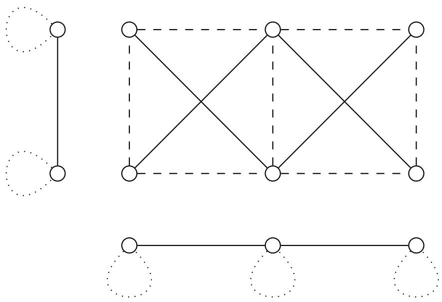

Figure 6.7: pseudo tensor product of graphs

the corresponding corner of $f^{i_1}_{\alpha_1,\overline{\alpha_2}}$ joins 1-paths $(e_{\alpha_1}, e_{\alpha_2})$ and $(e_{\alpha_1}, e_{\alpha_2'})$. We have similar result for a face corner at (v, u) generated by a face $f_{\overline{\alpha_1}} \in X_1$ and an edge $e_{\alpha_2} \in X_2$. By Definition 6.1, this tells that $L(X_1, v)\Box L(X_2, u) \cong L(X_1 \diamond X_2, (v, u))$. $\square$

Remark. Here is another viewpoint to understand the result. As we mentioned earlier, the zigzag product has faces generated by degenerate 2-gons and faces, with duplicated faces removed. In link graphs, it is like having a pseudo loop at each vertex, and performing a pseudo tensor product where a loop does not duplicate edges. This creates those dashed edges as illustrated in Figure 6.7, and is essentially the Cartesian product.

Now we investigate the polygonality of the zigzag product. Consider the complex X of two odd polygons glued together at one vertex. Note that $K_2 \diamond X$ has two faces, which meet at two non-adjacent vertices. A way to visualize this is to juxtapose two copies of Figure 6.4, and it is clear that when two pentagons meet at a vertex, the resulting two 10-gons above meet at two non-adjacent vertices. Hence the zigzag product does not preserve polygonality, and in fact it rarely produces polygonal complexes.

Proposition 6.5. Suppose that X and Y are connected polygonal complexes, X has at least two vertices, and Y has at least one face. Then the zigzag product $X \diamond Y$ is polygonal if and only if X is K_2, and Y has no faces of odd lengths meeting together.

Proof. We prove the only if part first. Assume that X has a vertex of valency higher than 1. Then we can find a path P of length 2 in X. Take an arbitrary face F in Y, and consider $P \diamond F$, which looks like stacking two polygonal columns in either Figure 6.3 or 6.4, depending on the parity of the length of F. We can find two faces generated by two different edges of P meeting each other every two zigzags, and hence $X \diamond Y$ is not polygonal, a contradiction. Therefore X has no vertex of valency higher than 1. By the connectedness and polygonality of X, X can only be K_2 without any face attached. And as explained in the paragraph above, Y has no faces of odd lengths meeting together.

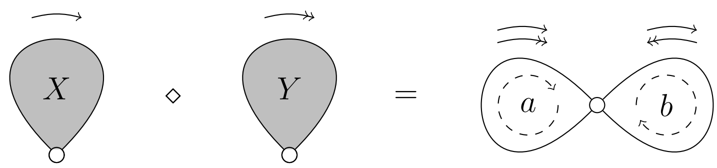

Figure 6.8: zigzag product of two 1-gons

Now we prove the if part. Since X^1 and Y^1 are simple graphs, the 1-skeleton of $X \diamond Y$, namely $X^1 \otimes Y^1$, is a simple graph. Note that faces in $X \diamond Y = K_2 \diamond Y$ are attached along simple closed cycles, as illustrated in Figures 6.3 and 6.4. Hence all we have to show is that the intersection of two faces in $K_2 \diamond Y$ is either empty, a vertex, or an edge. If two faces in $K_2 \diamond Y$ are generated by the same face F of Y, then F is of even length, and these two faces do not intersect, as illustrated in Figure 6.3. If two faces in $K_2 \diamond Y$ are generated by different faces of Y, since Y has no faces of odd lengths meeting together, by drawing polygonal columns as in Figure 6.3 or Figure 6.4 and juxtaposing them, it is easy to visualize that these two faces satisfy the condition. □

The zigzag product does not preserve simple connectedness either.

Proposition 6.6. Let X and Y be an n-gon and m-gon respectively, $1 \leq n \leq m$. Then $X \diamond Y$ is simply-connected if and only if $n = 1$ and m is even, or $n = 2$ and m is odd.

Proof. Note that X has n vertices, n edges, and 1 face, whereas Y has m vertices, m edges, and 1 face. By Definition 6.3, the complex $X \diamond Y$ has nm vertices, $2nm$ edges, and $m(n, 2) + n(m, 2)$ faces. When $n \geq 4$, $X \diamond Y$ has Euler characteristic

$$\chi(X \otimes Y) = nm - 2nm + m(n, 2) + n(m, 2) \leq -nm + 2n + 2m \leq 2n - 2m \leq 0.$$

When $n = 3$, $X \diamond Y$ has Euler characteristic

$$\chi(X \otimes Y) = nm - 2nm + m(n, 2) + n(m, 2) = -2m + 3(m, 2) \leq -2m + 6 \leq 0.$$

By Proposition 2.2, we know $X \diamond Y$ is not simply-connected when $n \geq 3$.

Consider the case when $n = m = 1$, namely the zigzag product of two 1-gons, as illustrated in Figure 6.8. The face of X and the edge of Y generate a face in $X \diamond Y$, attached along the single arrow, back and forth the double arrow, namely loop a followed by loop b. Similarly, the edge of X and the face of Y generate a face in $X \diamond Y$ attached along loop a followed by the reverse of loop b. Therefore $X \diamond Y$ has fundamental group $\langle\, a, b \mid ab = ab^{-1} = 1 \,\rangle \cong \mathbb{Z}_2$, and $X \diamond Y$ is not simply-connected.

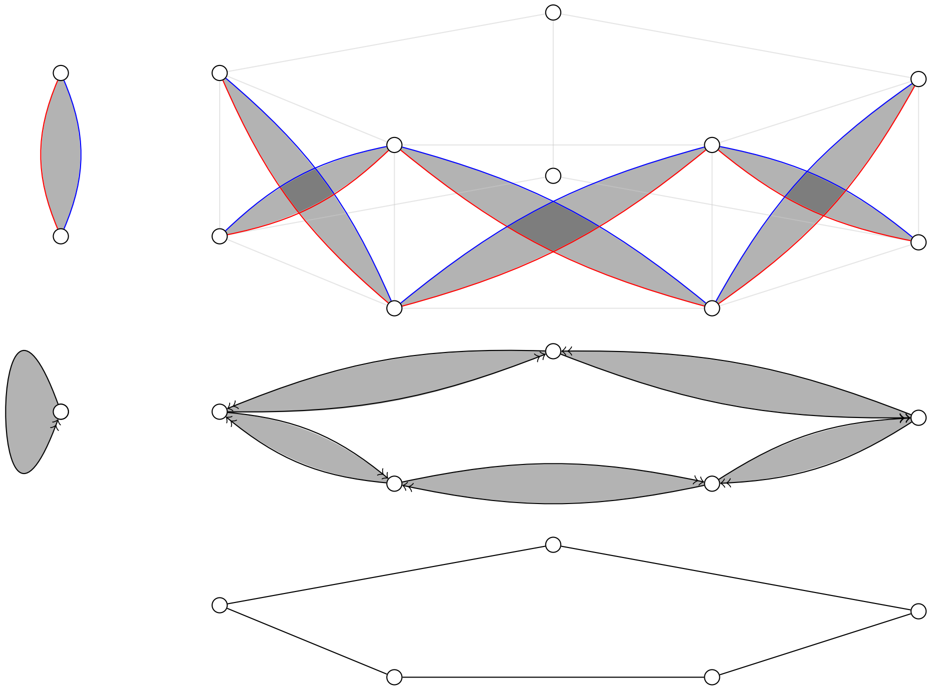

Figure 6.9: zigzag product of a 2-gon plus a 1-gon with a cycle

Suppose that $n = 1$ and $m \geq 2$. Note that the 1-gon of X and an edge of Y generate a 2-gon in $X \diamond Y$, attached along the orientation of the 1-gon in the X coordinate, back and forth the edge in the Y coordinate. Hence X and Y^1 generate a necklace of m 2-gons, as illustrated in Figure 6.9. When m is odd, the face generated by X^1 and Y has length $2m$, and wraps around the necklace of 2-gons twice. Note that $X \diamond Y$ is homotopic equivalent to a loop with a face wrapping around twice, which has fundamental group $\mathbb{Z}_2$, and hence $X \diamond Y$ is not simply-connected. When m is even, X^1 and Y generate two faces, each of which wraps around the necklace once. Note that these two face have no common edges, and $X \diamond Y$ is actually a sphere, which is of course simply-connected.

Now we work on the case of $n = 2$. Note that the 2-gon of X and an edge of Y generate two disjoint 2-gons in $X \diamond Y$. When m is odd, X and Y^1 generate a necklace of $2m$ 2-gons, as illustrated in Figure 6.9 with some 2-gons omitted. The face generated by Y and the blue (red) edge in X has length $2m$, and wraps around the necklace once along blue (red) edges. Hence $X \diamond Y$ is actually a sphere, and hence simply-connected. When m is even, X and Y^1 generate two necklaces of m 2-gons, and X^1 and Y generate 4 faces attached to these two necklaces. It is not hard to see that $X \diamond Y$ is the disjoint union of two spheres, and hence not simply-connected. In fact, since the 1-skeletons of X and Y are both bipartite, Theorem 3.4 implies that $X \diamond Y$ is not connected. $\square$

$$\begin{tikzcd}
 & \otimes\Gamma_i \arrow[d, "\pi_{\Gamma_j}"] \\
\Gamma \arrow[ur, dashed, "\exists!\psi"] \arrow[r, "\varphi_j"'] & \Gamma_j
\end{tikzcd}$$

Figure 6.10: universal property of $\otimes\Gamma_i$

In the proof of Proposition 6.6, we occasionally discover that for an odd number m, the zigzag product of a 2-gon with an m-gon is isomorphic to the zigzag product of a 1-gon with a $2m$-gon, a surprising result at first glance. Note that a 2-gon is the product of a 1-gon with K_2, and a $2m$-gon is the product of K_2 with an m-gon. Therefore we have

$$(\text{1-gon} \diamond K_2) \diamond m\text{-gon} = \text{1-gon} \diamond (K_2 \diamond m\text{-gon}),$$

which is less surprising if we know zigzag product is actually associative. For zigzag product to be associative, at least the graph tensor product of 1-skeletons needs to be associative, an immediate categorical result from the universal property in Proposition 5.4. Through the universal property, we can also generalize Corollary 5.6 easily.

Proposition 6.7. For $i \in \{1, \dots, n\}$, let Γ_i be a graph, and P_i be a path in Γ_i of length m from v_i to u_i. Then in $\otimes_{i=1}^n \Gamma_i$, there exists a unique path, denoted by $(P_1, \dots, P_n)_\otimes$, from $(v_1, \dots, v_n)$ to $(u_1, \dots, u_n)$ such that $\pi_{\Gamma_j}((P_1, \dots, P_n)_\otimes) = P_j$ for each $j \in \{1, \dots, n\}$.

Proof. By standard category theory, $\otimes_{i=1}^n \Gamma_i$ is the graph with projection maps $\pi_{\Gamma_j} : \otimes_{i=1}^n \Gamma_i \to \Gamma_j$ such that for any graph Γ and homomorphisms $\varphi_j : \Gamma \to \Gamma_j$, there exists a unique homomorphism $\psi : \Gamma \to \otimes_{i=1}^n \Gamma_i$ such that $\varphi_j = \pi_{\Gamma_j} \circ \psi$ for all j. In other word, there exists unique ψ such that the diagram in Figure 6.10 commutes for all j.

Let I be a graph which is a path of length m. Then there is a natural homomorphism φ_j from I to P_j for all $j \in \{1, \dots, n\}$. By the universal property above, there exists a unique homomorphism $\psi : I \to \otimes_{i=1}^n \Gamma_i$ such that $\varphi_j = \pi_{\Gamma_j} \circ \psi$ for all j. Hence we have $P_j = \varphi_j(I) = \pi_{\Gamma_j} \circ \psi(I)$ for all j. Note that $\psi(I)$ satisfies the condition for $(P_1, \dots, P_n)_\otimes$, and the uniqueness of $(P_1, \dots, P_n)_\otimes$ follows from the uniqueness of ψ. $\square$

Remark. The above proposition implies $((P_1, \dots, P_{n-1})_\otimes, P_n)_\otimes = (P_1, \dots, P_{n-1}, P_n)_\otimes$, and in fact we can bracket $P_1, \dots, P_n$ arbitrarily to get $(P_1, \dots, P_n)_\otimes$.

Proposition 6.8. For $j \in \{1,2,3\}$, let X_j be a polygonal cell complex with edge set $E(X_j) = \{e_{\alpha_j} \mid \alpha_j \in E_j\}$ and face set $F(X_j) = \{f_{\overline{\alpha_j}} \mid \overline{\alpha_j} \in F_j\}$. Let X be a polygonal cell complex with 1-skeleton $X_1 \otimes X_2 \otimes X_3$, and face set

$$F(X) = \left\{ f^{i_2,i_3}_{\overline{\alpha_1},\alpha_2,\alpha_3}, f^{i_1,i_3}_{\alpha_1,\overline{\alpha_2},\alpha_3}, f^{i_1,i_2}_{\alpha_1,\alpha_2,\overline{\alpha_3}} \mid \alpha_j \in E_j, \overline{\alpha_j} \in F_j \right\},$$

where $f^{i_2,i_3}_{\overline{\alpha_1},\alpha_2,\alpha_3}$ is a face of length $[2, n_{\overline{\alpha_1}}]$ attached along the length $n_{\overline{\alpha_1}}$ boundary of $f_{\overline{\alpha_1}}$ in the X_1 coordinate, zigzagging along e_{α_2} and e_{α_3} in the X_2 and the X_3 coordinates, with $i_2, i_3 \in \{0,1\}$ indicating the starting vertices of e_{α_2} and e_{α_3}, and we identify $f^{i_2,i_3}_{\overline{\alpha_1}\alpha_2,\alpha_3}$ with $f^{1-i_2,1-i_3}_{\overline{\alpha_1},\alpha_2,\alpha_3}$ when $n_{\overline{\alpha_1}}$ is odd. $f^{i_1,i_3}_{\alpha_1,\overline{\alpha_2},\alpha_3}$ and $f^{i_1,i_2}_{\alpha_1,\alpha_2,\overline{\alpha_3}}$ are similarly defined. Then we have

$$X \cong (X_1 \diamond X_2) \diamond X_3 \cong X_1 \diamond (X_2 \diamond X_3).$$

Proof. We only show $X \cong (X_1 \diamond X_2) \diamond X_3$ here, and the same argument also applies to $X \cong X_1 \diamond (X_2 \diamond X_3)$. By Definition 6.3, there are two types of faces in $(X_1 \diamond X_2) \diamond X_3$:

(1) Faces generated by a face, $f^{i_2}_{\overline{\alpha_1},\alpha_2}$ or $f^{i_1}_{\alpha_1,\overline{\alpha_2}}$, in $X_1 \diamond X_2$ and an edge e_{α_3} in X_3.

$f^{i_2}_{\overline{\alpha_1},\alpha_2}$: Let $n_{\overline{\alpha_1}}$ be the length of $f_{\overline{\alpha_1}}$. Then $f^{i_2}_{\overline{\alpha_1},\alpha_2}$ is of length $[2, n_{\overline{\alpha_1}}]$, attached along the boundary of $f_{\overline{\alpha_1}}$ in the X_1 coordinate, zigzagging along e_{α_2} in the X_2 coordinate, with $i_2 \in \{0, (2, n_{\overline{\alpha_1}}) - 1\}$. Since $f^{i_2}_{\overline{\alpha_1},\alpha_2}$ is of even length, $f^{i_2}_{\overline{\alpha_1},\alpha_2}$ and e_{α_3} generate 2 faces $(f^{i_2}_{\overline{\alpha_1},\alpha_2})^{i_3}_{\alpha_3}$ in $(X_1 \diamond X_2) \diamond X_3$, where $i_3 \in \{0,1\}$.

$f^{i_1}_{\alpha_1,\overline{\alpha_2}}$: Let $n_{\overline{\alpha_2}}$ be the length of $f_{\overline{\alpha_2}}$. Then $f^{i_1}_{\alpha_1,\overline{\alpha_2}}$ is of length $[2, n_{\overline{\alpha_2}}]$, attached along the boundary of $f_{\overline{\alpha_2}}$ in the X_2 coordinate, zigzagging along e_{α_1} in the X_1 coordinate, with $i_1 \in \{0, (2, n_{\overline{\alpha_2}}) - 1\}$. Since $f^{i_1}_{\alpha_1,\overline{\alpha_2}}$ is of even length, $f^{i_1}_{\alpha_1,\overline{\alpha_2}}$ and e_{α_3} generate 2 faces $(f^{i_1}_{\alpha_1,\overline{\alpha_2}})^{i_3}_{\alpha_3}$ in $(X_1 \diamond X_2) \diamond X_3$, where $i_3 \in \{0,1\}$.

(2) Faces generated by an edge $e^{\delta}_{\alpha_1,\alpha_2}$ in $X_1 \diamond X_2$ and a face $f_{\overline{\alpha_3}}$ of length $n_{\overline{\alpha_3}}$ in X_3. The edge $e^{\delta}_{\alpha_1,\alpha_2}$ is generated by e_{α_1} in X_1 and e_{α_2} in X_2, with $\delta \in \{0,1\}$ indicating whether e_{α_2} is flipped or not. Then $e^{\delta}_{\alpha_1,\alpha_2}$ and $f_{\overline{\alpha_3}}$ in X_3 generate 1 or 2 faces $f^{i}_{(\alpha_1,\alpha_2)^{\delta},\overline{\alpha_3}}$, where $i \in \{0, (2, n_{\overline{\alpha_3}}) - 1\}$ indicates the starting vertex of $e^{\delta}_{\alpha_1,\alpha_2}$.

Now we identify the 1-skeletons of X and $(X_1 \diamond X_2) \diamond X_3$. By Proposition 6.7, faces $f^{i_2,i_3}_{\overline{\alpha_1},\alpha_2,\alpha_3}$ and $(f^{i_2}_{\overline{\alpha_1},\alpha_2})^{i_3}_{\alpha_3}$, $f^{i_1,i_3}_{\alpha_1,\overline{\alpha_2},\alpha_3}$ and $(f^{i_1}_{\alpha_1,\overline{\alpha_2}})^{i_3}_{\alpha_3}$ as well, are attached along identical cycles in X and $(X_1 \diamond X_2) \diamond X_3$. Meanwhile, faces $f^{i_1,i_2}_{\alpha_1,\alpha_2,\overline{\alpha_3}}$ and $f^{i}_{(\alpha_1,\alpha_2)^{\delta},\overline{\alpha_3}}$ are attached along identical cycles in X and $(X_1 \diamond X_2) \diamond X_3$, where $i = i_1$ and $\delta = |i_1 - i_2|$. Note that when $n_{\overline{\alpha_1}}$ is odd, identifying $f^{i_2,i_3}_{\overline{\alpha_1}\alpha_2,\alpha_3}$ with $f^{1-i_2,1-i_3}_{\overline{\alpha_1},\alpha_2,\alpha_3}$ has the same effect as choosing i_2 from $\{0, (2, n_{\overline{\alpha_1}}) - 1\} = \{0, 0\}$. For odd $n_{\overline{\alpha_2}}$ or $n_{\overline{\alpha_3}}$ we have similar situations. Hence there is a bijection between faces of X and $(X_1 \diamond X_2) \diamond X_3$, and we have $X \cong (X_1 \diamond X_2) \diamond X_3$. $\square$

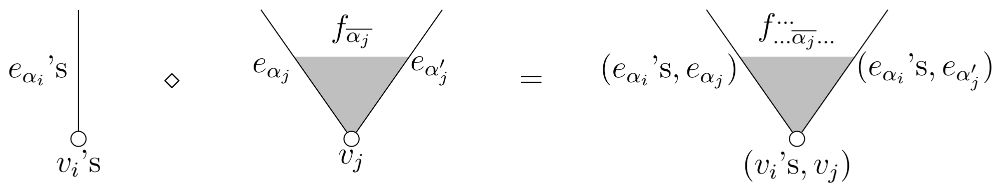

Figure 6.11: how a face corner is generated

Remark. The above description of X can easily be generalized to the zigzag product of n complexes $X_1 \diamond \ldots \diamond X_n$, where the first type of face of X is denoted $f^{i_2,\ldots,i_n}_{\overline{\alpha_1},\alpha_2,\ldots,\alpha_n}$. The key fact is that every face of X is generated by one face $f_{\overline{\alpha_j}}$ of X_j with one edge e_{α_i} of X_i from every $i \neq j$. Moreover, every face corner of X is generated by one face corner of X_j with one 1-path e_{α_i} of X_i from every $i \neq j$, as illustrated in Figure 6.11.

Theorem 6.9. Suppose that X is a platonic polygonal cell complex. Then the zigzag product of n copies of X is also a platonic complex.

Proof. In case X has no faces, then $\diamond_{i=1}^n X$ is the graph $\otimes_{i=1}^n X^1$, of which the platonicity follows easily from the definition of graph tensor product. Hereafter we assume that X has at least one face. Since X is platonic, X has the property that each vertex is incident to an edge, and each edge is incident to a face. By the remark above, this implies that $\diamond_{i=1}^n X$ has this property as well. Hence it suffices to show $\diamond_{i=1}^n X$ is flag-transitive.

Note that automorphisms of each X in the product and permutations of the n copies of X generate a subgroup of the automorphism group of $\diamond_{i=1}^n X$. In other words, we have

$$\mathrm{Aut}(X) \wr S_n \leq \mathrm{Aut}(\diamond_{i=1}^n X).$$

For any two face corners A and B of $\diamond_{i=1}^n X$, consider the generating 1-paths and face corners from each factor. We can find $\rho \in \mathrm{Aut}(X) \wr S_n$ sending each generating 1-path of A to a generating 1-path of B, and the generating face corner of A to the generating face corner of B. Then such ρ is actually an automorphism of $\diamond_{i=1}^n X$ sending A to B. This shows that $\mathrm{Aut}(\diamond_{i=1}^n X)$ acts transitively on face corners with orientations, and therefore transitively on half-corners. In other words, $\mathrm{Aut}(\diamond_{i=1}^n X)$ acts transitively on fev flags, and the zigzag product $\diamond_{i=1}^n X$ is a platonic complex. $\square$

Remark. Let G_{fev} and G_{ev} be the fev and ev stabilizers of X respectively. Then in fact there are $(n-1)! \cdot |G_{ev}|^{n-1} \cdot |G_{fev}|$ such $\rho \in \mathrm{Aut}(X) \wr S_n$ sending A to B.

Now we can construct platonic complexes with Q_n links through zigzag product.

Corollary 6.10. Let X be a polygon. Then $\diamond_{i=1}^{n} X$ is a platonic complex with Q_n links.

Proof. Note that the polygon X is a platonic complex with K_2 links. By Theorem 6.9, the zigzag product $\diamond_{i=1}^{n} X$ is a platonic complex. By Theorem 6.4, each vertex in $\diamond_{i=1}^{n} X$ has link graph $\square_{i=1}^{n} K_2 = Q_n$. This completes the proof. □

To end this chapter, we use zigzag product to give more platonic complexes.

Proposition 6.11. Let Γ be a symmetric graph, and X be a platonic polygonal complex. Then the zigzag product $\Gamma \diamond X$ is also a platonic complex.

Proof. The proof is basically identical to the proof of Theorem 6.9. □

Chapter 7

Tensor Symmetry

At the end of Chapter 5, we raised a question about the automorphism group of the complex tensor product: for any two complexes X and Y, is the following relation true?

$$\mathrm{Aut}(X \otimes Y) \cong \mathrm{Aut}(X) \times \mathrm{Aut}(Y)$$

When X an Y are isomorphic, we can swap X and Y to obtain an extra automorphism, since the complex tensor product is commutative up to isomorphism. In addition to swapping, the following proposition gives more automorphisms in a less obvious way.

Proposition 7.1. Let X, Y, and Z be polygonal cell complexes. Then we have

$$(X \otimes Y) \otimes Z \cong X \otimes (Y \otimes Z).$$

In other words, complex tensor product is associative up to isomorphism.

Proof. A categorical result of the universal property in Proposition 5.12 $\square$

The associativity of the complex tensor product complicates $\mathrm{Aut}(X \otimes Y)$. For example, if Y can be factorized into $X \otimes Z$, then $X \otimes Y \cong X \otimes (X \otimes Z)$ has an automorphism swapping the two copies of X. Hence the symmetry of the product of complexes is also related to the factoring of complexes. In response to associativity, we modify the original question as follows: for complexes X_i which are irreducible with respect to complex tensor product, is the automorphism group $\mathrm{Aut}(\otimes X_i)$ generated by automorphisms of X_i's, together with permutations of isomorphic factors? By a **Cartesian automorphism**, we mean an element in the subgroup of $\mathrm{Aut}(\otimes X_i)$ generated in the above manner.

There have been lots of studies about the symmetry of different products of graphs. One of the major goals of this chapter is to apply the theory of the graph direct product to the complex tensor product. Hence we first introduce related theorems about the graph

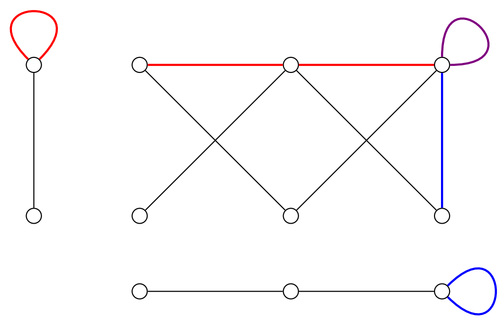

Figure 7.1: direct product of graphs in $\mathfrak{S}_0$

direct product. The book [15] by Hammack, Imrich, and Klavžar offers a comprehensive survey of products of graphs, and we shall follow their approach and terminology here.

We have defined the direct product of graphs in Chapter 3. Here we give the definition again, with an emphasis on the possible presence of loops. We say that a graph Γ is a **simple graph with loops admitted** if for any $u, v \in V(\Gamma)$, there is at most one edge joining u and v, including the case $u = v$. In particular, there is at most one loop at a vertex. For convenience, we use $\mathfrak{S}$ to denote the class of simple graphs, and $\mathfrak{S}_0$ to denote the class of simple graphs with loops admitted.

Definition 7.2. Let Γ and Γ' be two graphs in $\mathfrak{S}_0$. The **direct product** of Γ and Γ', denoted by $\Gamma \times \Gamma'$, is a graph in $\mathfrak{S}_0$ with vertex set $V(\Gamma \times \Gamma') = V(\Gamma) \times V(\Gamma')$. There is an edge joining two vertices (v, v') and (u, u') in $\Gamma \times \Gamma'$ if and only if there is an edge joining v and u in Γ, and there is an edge joining v' and u' in Γ'.

Note in the above definition, v and v' could be the same vertex, as well as u and u'. Figure 7.1 illustrates the direct product of two graphs in $\mathfrak{S}_0$. Under this definition, notice that a loop L serves as the identity element of direct product of graphs. In other words, for any simple graph Γ with loops admitted, we always have

$$L \times \Gamma \cong \Gamma \times L \cong \Gamma.$$

A graph Γ is **prime** if Γ has more than one vertex, and $\Gamma \cong \Gamma_1 \times \Gamma_2$ implies that either Γ_1 or Γ_2 is a loop. Note that the idea of being prime depends on the class of graphs we are talking about. For example, let Γ be a path of length 3, which has 4 vertices. Then Γ is prime in $\mathfrak{S}$, as the only possible factoring is the product of two edges, which is the disjoint union of two edges. And the statement that $\Gamma \cong \Gamma_1 \times \Gamma_2$ implies either Γ_1 or Γ_2 is a loop is still logically true. However, Γ can be factorized in $\mathfrak{S}_0$ as the graph on the left of Figure 7.1 times one edge in the bottom, and hence Γ is not prime in $\mathfrak{S}_0$.

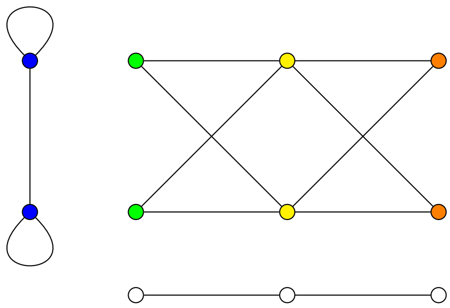

Figure 7.2: vertices with the same set of neighbours

Consider the question of factoring a graph into the product of prime graphs. For a finite graph, such a prime factorization always exists, since the number of vertices of factors decreases as the factoring goes. However, such a prime factorization is not necessarily unique, and it depends on the graph itself and the class of graphs where we do the factoring. For example, a path of length 3 together with associativity can be used to create graphs with non-unique prime factorizations in $\mathfrak{G}$. There are also graphs with non-unique prime factorizations in $\mathfrak{G}_0$, an example of which can be found in [15]. The following theorem of unique prime factorization is due to McKenzie [24].

Theorem 7.3. Suppose that $\Gamma \in \mathfrak{G}_0$ is a finite connected non-bipartite graph with more than one vertex. Then Γ has a unique factorization into primes in $\mathfrak{G}_0$.

The next question is about the automorphism group of direct product, which hopefully has only these Cartesian automorphisms with respect to the product. Note that a pair of vertices with the same set of neighbours creates pairs of vertices with the same set of neighbours in the direct product, and results in lots of non-Cartesian automorphisms. This phenomenon is illustrated in Figure 7.2 , where a vertex with a loop should have itself as a neighbour. We say that a graph is ***R*-thin** if there are no vertices with the same set of neighbours. In addition to R-thinness, the disconnectedness due to Theorem 3.4 also creates non-Cartesian automorphisms. Even when the direct product is connected, the example of $K_2 \times$ dodecahedron graph in Figure 3.9 still has an exotic automorphism. The following theorem is due to Dörfler [10].

Theorem 7.4. Suppose that $\Gamma \in \mathfrak{G}_0$ is a finite connected non-bipartite R-thin graph with a prime factorization $\Gamma = \Gamma_1 \times \Gamma_2 \times \cdots \times \Gamma_n$ in $\mathfrak{G}_0$. Then $\mathrm{Aut}(\Gamma)$ is generated by automorphisms of prime factors and permutations of isomorphic factors.

We would like to use Theorems 7.3 and 7.4 to develop similar results for the complex tensor product. The first problem we immediately encounter is that, for the complex tensor product, we obtain the 1-skeleton of the product through the graph tensor product, which is not exactly the same as the direct product of graphs. Fortunately, such a difference does not really take place in the context of this thesis.

Proposition 7.5. Let $\Gamma \in \mathfrak{S}_o$ be a finite connected non-bipartite R-thin graph with more than one vertex, and $\Gamma = \Gamma_1 \times \Gamma_2 \times \cdots \times \Gamma_n$ be the unique prime factorization in $\mathfrak{S}_0$. If Γ is edge-transitive, then Γ and each prime factor Γ_i are in $\mathfrak{S}$.

Proof. Since Γ has more than one vertex, the connectedness of Γ implies that Γ has a non-loop edge. By the edge-transitivity of Γ, we know Γ has no loop, and hence is in $\mathfrak{S}$. If each factor Γ_i has a loop, then the product Γ will have a loop, which is not true. If each factor Γ_i is loop-free, then we have finished the proof. Hence we can assume there is at least one factor with a loop, and at least one factor without a loop.

Let Γ_α be the direct product of all factors with a loop, and Γ_β be the direct product of all factors without a loop. Then we have $\Gamma = \Gamma_\alpha \times \Gamma_\beta$. Note that permuting isomorphic factors of Γ does not involve permuting factors of Γ_α with factors of Γ_β. By Theorem 7.4, we have $\mathrm{Aut}(\Gamma) = \mathrm{Aut}(\Gamma_\alpha) \times \mathrm{Aut}(\Gamma_\beta)$. Since a prime factor has more than one vertex, Γ_α and Γ_β both have more than one vertex. Since Γ is connected, Γ_α and Γ_β are both connected. Hence Γ_α has a loop at some vertex v and a non-loop edge joining two vertices v_α and v'_α, while Γ_β has a non loop edge joining two vertices v_β and v'_β. Then in $\Gamma = \Gamma_\alpha \times \Gamma_\beta$, there is an edge joining (v, v_β) and (v, v'_β), and another edge joining (v_α, v_β) and (v'_α, v'_β). Notice that $\mathrm{Aut}(\Gamma) = \mathrm{Aut}(\Gamma_\alpha) \times \mathrm{Aut}(\Gamma_\beta)$ can not send the first edge to the second one, contradicting the assumption that Γ is edge-transitive. □

Remark. To visually interpret the last few lines of the proof, it says that a Cartesian automorphism can not permute horizontal edges with slant edges in Figure 7.2.

Now we move on to the factorization of polygonal cell complexes. First consider the following example. Let X and Y be a triangle and a pentagon respectively, X' be a cycle of length 3 with two triangles attached, and Y' be a cycle of length 5 with two pentagons attached. Since the numbers of vertices of these complexes are prime, the only possible way to factorize them is to have a factor of one vertex with at least a loop and a face, which creates double edges in the product. Hence we know these complexes can not be factorized further, and we have non-unique factorizations $X \otimes Y' \cong X' \otimes Y$.

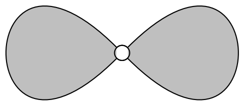

Figure 7.3: a simple complex with two 1-gons

Here we give another example of non-unique factorization. Let X be a triangle, and Y' be a $(7 \cdot 5)$-gon wrapped around a cycle of length 5. By Definition 5.7, since 3 and $7 \cdot 5$ are coprime, $X \otimes Y'$ has two faces of length $3 \cdot 5 \cdot 7$, wrapped around two cycles of length $3 \cdot 5$ for 7 rounds. Consider a $(7 \cdot 3)$-gon X' wrapped around a cycle of length 3, and a pentagon Y. It is easy to see that $X \otimes Y' \cong X' \otimes Y$, and these complexes can not be factorized further. To avoid these non-uniquely factorized situations, we restrict our discussion to the factorization of simple complexes.

Definition 7.6. A polygonal cell complex X is a **simple** complex if X has at least one face, X has no pairs of faces attached along the same cycle, and the attaching map of each face does not wrap around a cycle more than once. A polygonal cell complex X is a **prime** complex if there do not exist complexes X_1 and X_2 such that $X = X_1 \otimes X_2$.

Remark. Figure 7.3 above is a simple complex with two 1-gons. If we add another 2-gon attached along two different loops, the resulting complex is still a simple complex, as the boundary cycles of theses faces are not exactly the same.

To factorize a complex X, our general setting is as follows. We assume that we know a factorization of the 1-skeleton $X^1 = \Gamma_1 \otimes \Gamma_2$, and try to find a complex factorization $X = X_1 \otimes X_2$ such that $X_1^1 = \Gamma_1$ and $X_2^1 = \Gamma_2$. A natural thought is to project the faces of X down to Γ_1 and Γ_2 to be faces. Consider the complex tensor product of a triangle and a pentagon, which is a complex with two 15-gons. Note that when we project these two 15-gons back to the 1-skeletons of factors, what we obtain are 15-gons wrapped around cycles of length 3 and 5 respectively, not the original faces.

Definition 7.7. Let X be a polygonal cell complex, f be a face of X attached along a cycle C_f, and Γ_1 and Γ_2 be two graphs such that $X^1 = \Gamma_1 \otimes \Gamma_2$. The **reductive projection** of f to Γ_i, denoted by $\pi_{\Gamma_i}(f)$, is a face attached along the reduced cycle of $\pi_{\Gamma_i}(C_f)$ in Γ_i, namely the shortest cycle C such that repeating C gives $\pi_{\Gamma_i}(C_f)$.

Remark. In exactly the same way, we can define $\pi_{\Gamma_i}(f)$ for the case $X^1 = \otimes_{i=1}^{n} \Gamma_i$. Note that when $X^1 = \Gamma_1 \otimes \Gamma_2 \otimes \Gamma_3$, we have $\pi_{\Gamma_1}(f) = \pi_{\Gamma_1}(\pi_{\Gamma_1 \otimes \Gamma_3}(f)) = \pi_{\Gamma_1}(\pi_{\Gamma_1 \otimes \Gamma_2}(f))$.

Proposition 7.8. Let X be a simple complex, and Γ_1 and Γ_2 be two graphs such that $X^1 = \Gamma_1 \otimes \Gamma_2$. If there exist two complexes X_1 and X_2 with 1-skeletons Γ_1 and Γ_2 respectively such that $X = X_1 \otimes X_2$, then X_1 and X_2 are simple complexes whose faces are precisely the reductive projections of faces of X.

Proof. Suppose that such complexes X_1 and X_2 exist. Let f be a face of X attached along a cycle C_f of length n, and let C_j of length n_j be the reduced cycle of $\pi_{\Gamma_j}(C_f)$ in X_j for $j \in \{1,2\}$. Note that f is generated by a face f_1 of X_1 attached along m_1C_1, and by a face f_2 of X_2 attached along m_2C_2, where m_iC_i is the cycle made by repeating C_i for m_i times. By Definition 5.7, f_1 and f_2 generate faces attached along $(m_1C_1, m_2C_2)_{\otimes}^{i^\delta}$, where $i \in \{0, 1, \ldots, (m_1n_1, m_2n_2) - 1\}$ and $\delta \in \{0, 1\}$. By the Euclidean algorithm, we can find an integer $k > 0$ such that $k \equiv 0 \mod n_1$ and $k \equiv (n_1, n_2) \mod n_2$. Note that in k steps along $(m_1C_1, m_2C_2)_{\otimes}^{0^\delta}$, we can walk from the starting vertex of $(m_1C_1, m_2C_2)_{\otimes}^{0^\delta}$ to the starting vertex of $(m_1C_1, m_2C_2)_{\otimes}^{(n_1,n_2)^\delta}$, so these two cycles are identical. Since X is simple, there are no pairs of faces attached along the same cycle in X. Therefore we have $(n_1, n_2) \geq (m_1n_1, m_2n_2) \geq (n_1, n_2)$. Now consider the length of the face f, which is

$$n = [m_1n_1, m_2n_2] = \frac{m_1n_1 \cdot m_2n_2}{(m_1n_1, m_2n_2)} = \frac{m_1m_2 \cdot n_1n_2}{(n_1, n_2)} = m_1m_2 \cdot [n_1, n_2].$$

This shows that f is attached along some cycle $(C_1, C_2)_{\otimes}^{i^\delta}$ of length $[n_1, n_2]$ for m_1m_2 rounds, and the simplicity of X implies that $m_1 = m_2 = 1$. In other words, X_i must have the reductive projection $\pi_{\Gamma_i}(f)$ of f as its face. Note that different faces of X might have the same reductive projection in X_i, and we have to discard duplicated ones. Otherwise duplicated faces in X_i will generate duplicated faces in X, violating the simplicity of X. Conversely, any faces f_1 of X_1 and f_2 of X_2 are the reductive projections of the faces in X they generate. Hence X_1 and X_2 are the simple complexes with exactly those faces from the reductive projections of faces of X. $\square$

Proposition 7.9. Let X, X_1, and X_2 be polygonal cell complexes such that $X = X_1 \otimes X_2$. Then X is a simple complex if and only if X_1 and X_2 are simple complexes.

Proof. Proposition 7.8 takes care of the only if part, and here we prove the if part. Suppose that X has an n-gon f attached along a cycle for m rounds. Since X_1 and X_2 are simple, f must be generated by the reductive projections of f to X_1^1 and X_2^1, which are of length l_1 and l_2 respectively. Note that l_1 and l_2 both divide $\frac{n}{m}$. Then the two reductive projections generate faces of length $n = [l_1, l_2] \leq \frac{n}{m}$. Hence we can conclude that $m = 1$. If there is another face f' in X attached along the same cycle with f, then f' is also generated by the reductive projections of f. If we can show a face in X_1 and a face in X_2 do not generate duplicated faces in X, then this implies X is a simple complex.

Suppose that a face f_1 of X_1 has vertices $v_0, v_1, \ldots, v_{p-1}, v_0$ in order, and a face f_2 of X_2 has vertices $u_0, u_1, \ldots, u_{q-1}, u_0$ in order. By the remark after Definition 5.7, every pair of corners of f_1 and f_2 appears exactly once in the faces generated by f_1 and f_2. If two faces generated by f_1 and f_2 are attached along the same cycle in X, there must be two pairs of corners of f_1 and f_2 forming the same corner in X. In particular, we can find $(v_i, u_{i'}) = (v_j, u_{j'})$ such that $i \neq j$ or $i' \neq j'$. When $i \neq j$, we have $v_i = v_j$ and $v_{i+k} = v_{j+k}$ for any integer k mod p. This implies that f_1 wraps around a cycle more than once, violating the simplicity of X_1. Similarly $i' \neq j'$ contradicts the simplicity of X_2. The contradiction results from the assumption that two faces generated by f_1 and f_2 are attached along the same cycle in X. Hence we know that f_1 and f_2 does not generate duplicated faces, and the simplicity of X follows. □

Proposition 7.10. Let X be a simple complex, and Γ_1 and Γ_2 be two graphs such that $X^1 = \Gamma_1 \otimes \Gamma_2$. Then the following two statements are equivalent:

(1) There exist two complexes X_1 and X_2 such that $X_i^1 = \Gamma_i$ and $X = X_1 \otimes X_2$.

(2) For any faces f_1 and f_2 of X, X contains all faces generated by $\pi_{\Gamma_1}(f_1)$ and $\pi_{\Gamma_2}(f_2)$.

Proof. Assume (1). By Proposition 7.8, X_1 and X_2 are the simple complexes with exactly those reductive projections of X as faces. For any faces f_1 and f_2 of X, $\pi_{\Gamma_1}(f_1)$ is a face of X_1, and $\pi_{\Gamma_2}(f_2)$ is a face of X_2. Since $X = X_1 \otimes X_2$, X contains all faces generated by $\pi_{\Gamma_1}(f_1)$ and $\pi_{\Gamma_2}(f_2)$. Hence (1) implies (2).

Assume (2). First we show that a face f of X can be generated by $\pi_{\Gamma_1}(f)$ and $\pi_{\Gamma_2}(f)$. Let C_f, C_1, and C_2 be the boundary cycles of f, $\pi_{\Gamma_1}(f)$, and $\pi_{\Gamma_2}(f)$ respectively. By Definition 7.7, we can assume that $\pi_{\Gamma_j}(C_f) = n_j C_j$ for $j \in \{1, 2\}$, namely repeating C_j for n_j times gives $\pi_{\Gamma_j}(C_f)$. Note that f is attached along some cycle $(n_1 C_1, n_2 C_2)_{\otimes}^{i^\delta}$, which can be rewritten as $(n_1, n_2)(\frac{n_1}{(n_1,n_2)} C_1, \frac{n_2}{(n_1,n_2)} C_2)_{\otimes}^{i^\delta}$. Since the simple complex X has no face attached around a cycle more than once, we know that $(n_1, n_2) = 1$, and therefore

$$\text{length } C_f = n_1 \cdot (\text{length } C_1) = n_2 \cdot (\text{length } C_2) = [\text{length } C_1, \text{length } C_2].$$

This shows that $\pi_{\Gamma_1}(f)$ and $\pi_{\Gamma_2}(f)$ can generate the face f. Now let X_1 and X_2 be the simple complexes with exactly those faces from the reductive projections of X. By

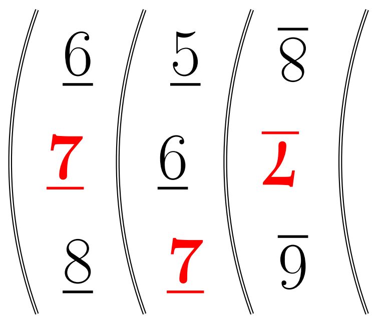

Figure 7.4: a face generated by 3 faces in complex tensor product

Proposition 7.9, $X_1 \otimes X_2$ is a simple complex, and in particular $X_1 \otimes X_2$ has no duplicated faces. By the assumption of (2), X contains all the faces of $X_1 \otimes X_2$. Conversely, any face f of X is a face of $X_1 \otimes X_2$, since f can be generated by $\pi_{\Gamma_1}(f)$ and $\pi_{\Gamma_2}(f)$. Then we have $X = X_1 \otimes X_2$, and hence (2) implies (1). $\square$

Although we already know the associativity of complex tensor product through the universal property, it will be helpful to understand how faces are formed in the product of more than two complexes. First let us review the product of two complexes. Let f_α be a face of length n_α attached along a cycle C_α in X, and f_β be a face of length n_β attached along a cycle C_β in Y. By Definition 5.7, f_α and f_β generate faces $f_{\alpha,\beta}^{i^\delta}$ of length $[n_\alpha, n_\beta]$ attached along $(C_\alpha, C_\beta)_\otimes^{i^\delta}$, $i \in \{0, 1, \ldots, (n_\alpha, n_\beta) - 1\}$, $\delta \in \{0, 1\}$. To explain the boundary cycle of $f_{\alpha,\beta}^{i^\delta}$ in plain language, basically we pick a pair of corners of f_α and f_β to start, and go around C_α and C_β in two coordinates respectively until we return to the starting pair of corners. Note that the index i is chosen in such a way that each pair of corners appears exactly once among all faces generated by f_α and f_β.

A good way to visualize this is a slot machine of two reels of length $[n_\alpha, n_\beta]$, cyclically labeled by the vertices of f_α and f_β respectively. Faces generated by f_α and f_β have a one-to-one correspondence with different combinations of two reels, with flipping allowed for the second reel. From this aspect, it is easy to see that for face f_j of length n_j in complex X_j, $j \in \{1, 2, \ldots, m\}$, $f_1, f_2, \ldots, f_m$ generate faces in $\otimes_{j=1}^m X_j$ of length $[n_1, n_2, \ldots, n_m]$ such that each m-tuple of corners appears exactly once among all generated faces. Faces generated by $f_1, f_2, \ldots, f_m$ have a one-to-one correspondence with different combinations of m reels of length $[n_1, n_2, \ldots, n_m]$, cyclically labeled by the vertices of f_j respectively, with flipping allowed from the second reel on. Figure 7.4 illustrates how a face is generated by the complex tensor product of 3 faces from such an aspect.

Theorem 7.11. Let X be a simple polygonal cell complex. If the 1-skeleton of X is a finite simple connected non-bipartite R-thin edge-transitive graph with more than one vertex, then X has a unique factorization into prime complexes.

Proof. By Theorem 7.3, since $X^1 \in \mathfrak{S} \subset \mathfrak{S}_0$ is a finite connected non-bipartite graph with more than one vertex, X^1 has a unique factorization $X^1 = \Gamma_1 \times \Gamma_2 \times \cdots \times \Gamma_n$ into primes in $\mathfrak{S}_0$ with respect to direct product of graphs. By Proposition 7.5, the edge-transitivity of X^1 implies that each prime factor Γ_i is in fact a simple graph. On the other hand, if we factorize X^1 with respect to graph tensor product, each factor would also be a simple graph with more than one vertex, because a loop creates double edges in the product, and a single vertex breaks the connectivity of the product. Note that direct product and graph tensor product coincide in $\mathfrak{S}$. Hence we know X^1 has a unique factorization $X^1 = \Gamma_1 \otimes \Gamma_2 \otimes \cdots \otimes \Gamma_n$ into primes in $\mathfrak{S}$ with respect to graph tensor product.

Now we consider the factorization of the complex X. Note that we can always obtain a prime factorization of X, since the number of vertices of factors decreases as the factoring goes. Suppose X has two factorizations A and B, and X_0 is a prime factor of X in A with 1-skeleton $\Gamma_1 \otimes \Gamma_2$. By Proposition 7.10, there exist two faces f_1 and f_2 such that X_0 lacks certain face generated by $\pi_{\Gamma_1}(f_1)$ and $\pi_{\Gamma_2}(f_2)$. In other words, there is certain pair of corners of $\pi_{\Gamma_1}(f_1)$ and $\pi_{\Gamma_2}(f_2)$ missing in the faces of X_0, and hence such pair will be absent in the n-tuples representing face corners of X. By Proposition 7.8, we can find faces $\overline{f_1}$ and $\overline{f_2}$ of X such that $\pi_{\Gamma_1 \otimes \Gamma_2}(\overline{f_1}) = f_1$ and $\pi_{\Gamma_1 \otimes \Gamma_2}(\overline{f_2}) = f_2$, and we have $\pi_{\Gamma_1}(\overline{f_1}) = \pi_{\Gamma_1}(f_1)$ and $\pi_{\Gamma_2}(\overline{f_2}) = \pi_{\Gamma_2}(f_2)$. If Γ_1 and Γ_2 belong to different prime factors X_1 and X_2 in B, we can reductively project $\overline{f_i}$ to X_i to obtain a face f_i' of X_i, $i \in \{1, 2\}$. Then we have $\pi_{\Gamma_1}(f_1') = \pi_{\Gamma_1}(\overline{f_1}) = \pi_{\Gamma_1}(f_1)$ and $\pi_{\Gamma_2}(f_2') = \pi_{\Gamma_2}(\overline{f_2}) = \pi_{\Gamma_2}(f_2)$. Notice that f_1' and f_2' generate all possible pairs of corners of $\pi_{\Gamma_1}(f_1)$ and $\pi_{\Gamma_2}(f_2)$ in $X_1 \otimes X_2$ and hence in X, a contradiction. So Γ_1 and Γ_2 belong to the same prime factor in B.

The above argument can be applied to the case when the 1-skeleton of X_0 is the graph tensor product of more than two prime graphs, simply by splitting prime graph factors into two groups. It follows that every prime 1-skeleton factor of X_0 belongs to the same prime complex X_0' in B. Conversely, every prime 1-skeleton factor of X_0' belongs to X_0, and hence X_0 and X_0' are actually the same. In case X_0 has a prime 1-skeleton Γ_j, then Γ_j belongs to some X_0' in B with a prime 1-skeleton, otherwise the prime 1-skeleton factors of X_0' belong to at least two complexes in A. In conclusion, we know two factorizations A and B are identical, and X has a unique factorization into prime complexes. $\square$

Theorem 7.12. Suppose that X is a simple polygonal cell complex, and its 1-skeleton is a finite simple connected non-bipartite edge-transitive R-thin graph with more than one vertex. Let $X = X_1 \otimes X_2 \otimes \cdots \otimes X_n$ be a prime factorization of X. Then $\mathrm{Aut}(X)$ is generated by automorphisms of prime factors and permutations of isomorphic factors.

Proof. Since X has no faces attached along the same cycle, an automorphism of X is completely determined by its action on the 1-skeleton X^1, and we can identify $\mathrm{Aut}(X)$ as a subgroup of $\mathrm{Aut}(X^1)$. To understand $\mathrm{Aut}(X^1)$, by the argument in the proof of Theorem 7.11, we know X^1 has a unique factorization $X^1 = \Gamma_1 \times \Gamma_2 \times \cdots \times \Gamma_m = \Gamma_1 \otimes \Gamma_2 \otimes \cdots \otimes \Gamma_m$ into primes in $\mathfrak{S}$. By Theorem 7.4, the extra R-thin condition on X^1 implies that $\mathrm{Aut}(X^1)$ is generated by automorphisms of Γ_i's and permutations of isomorphic Γ_j's.

Let φ be an arbitrary automorphism of X, which can be represented as some $\rho \in \times_{i=1}^{m} \mathrm{Aut}(\Gamma_i)$ followed by a permutation of Γ_j's. This implies that for any face f of X

$$\varphi\big(\pi_{\otimes_{i\in I}\Gamma_i}(f)\big) = \pi_{\varphi(\otimes_{i\in I}\Gamma_i)}\big(\varphi(f)\big) = \pi_{\otimes_{i\in I}\varphi(\Gamma_i)}\big(\varphi(f)\big),$$

where I is an arbitrary non-empty subset of $\{1, 2, \ldots, m\}$. Suppose that X_1 has 1-skeleton $X_1^1 = \otimes_{i\in I}\Gamma_i$ for some $I \subset \{1, 2, \ldots, m\}$. We claim that $\forall i \in I$, $\varphi(\Gamma_i)$ belongs to the same prime factor X_k of X. If not, then we can find $I_1 \sqcup I_2 = I$ such that $\forall i \in I_1, \forall j \in I_2$, $\varphi(\Gamma_i)$ and $\varphi(\Gamma_j)$ belong to different prime factors of X. Let $\Gamma_\alpha = \otimes_{i\in I_1}\Gamma_i$ and $\Gamma_\beta = \otimes_{j\in I_2}\Gamma_j$, and hence we have $X_1^1 = \Gamma_\alpha \otimes \Gamma_\beta$. Since X_1 is prime, by Proposition 7.10, we can find faces f_1 and f_2 of X_1 such that X_1 lacks certain face generated by $\pi_{\Gamma_\alpha}(f_1)$ and $\pi_{\Gamma_\beta}(f_2)$. By Proposition 7.8, we can find faces $\overline{f_1}$ and $\overline{f_2}$ of X such that $\pi_{\Gamma_\alpha\otimes\Gamma_\beta}(\overline{f_1}) = f_1$ and $\pi_{\Gamma_\alpha\otimes\Gamma_\beta}(\overline{f_2}) = f_2$. Then the complex X lacks certain corner combination of $\pi_{\Gamma_\alpha}(\overline{f_1})$ and $\pi_{\Gamma_\beta}(\overline{f_2})$ in the m-tuples representing face corners of X. By taking the automorphism φ, the complex X lacks certain corner combination of $\pi_{\otimes_{i\in I_1}\varphi(\Gamma_i)}(\varphi(\overline{f_1}))$ and $\pi_{\otimes_{j\in I_2}\varphi(\Gamma_j)}(\varphi(\overline{f_2}))$, which is impossible because $\varphi(\Gamma_i)$ and $\varphi(\Gamma_j)$ belong to different prime factors of X, and taking complex tensor product of these factors generates all the corner combinations.

Hence for every 1-skeleton factor Γ_i of X_1, $\varphi(\Gamma_i)$ belongs to the same prime factor X_k of X. By considering φ^{-1}, we know that X_k has exactly these $\varphi(\Gamma_i)$'s as 1-skeleton factors. Moreover, $\varphi\big(\pi_{\otimes_{i\in I}\Gamma_i}(f)\big) = \pi_{\otimes_{i\in I}\varphi(\Gamma_i)}\big(\varphi(f)\big)$ implies that φ induces an isomorphism from X_1 to X_k. This shows that every $\varphi \in \mathrm{Aut}(X)$ can be represented as some $\sigma \in \times_{i=1}^{n} \mathrm{Aut}(X_i)$ followed by a permutation of X_j's, and the theorem holds. □

Remark. Let $\widetilde{X}$ be the disjoint union of prime factors of X. Then the above theorem implies that $\mathrm{Aut}(X) \cong \mathrm{Aut}(\widetilde{X})$, which is a convenient way to describe $\mathrm{Aut}(X)$.

The following corollary is a partial converse of Theorem 5.15.

Corollary 7.13. Suppose that X is a simple polygonal cell complex, and its 1-skeleton is a finite simple connected non-bipartite edge-transitive R-thin graph with more than one vertex. If X is platonic, then any factor of X is platonic.

Proof. Note that it suffices to show that any prime factor of X is platonic. Then by Theorem 7.11 and Theorem 5.15, any factor of X is a complex tensor product of platonic prime factors of X, and hence is platonic. Since X is a simple complex, X has at least one face. The platonicity of X implies that each vertex of X is incident to an edge, and each edge of X is incident to a face. By considering the projection map, we know that any prime factor of X has this incidence property as well, and the platonicity of a prime factor of X is equivalent to being flag-transitive.

By Theorem 7.11, X has a unique prime factorization $X = X_1 \otimes X_2 \otimes \cdots \otimes X_n$. Suppose that one of the prime factors is not platonic, without loss of generality say X_1, and X_i is isomorphic to X_1 if and only of $1 \leq i \leq m$ for some integer $m \leq n$. Since X_1 is not platonic, there exist two oriented face corners (e_1^1, v_1, e_1^2) and $(e_1^{1'}, v_1', e_1^{2'})$ in X_1 such that $\mathrm{Aut}(X_1)$ can not map one corner to the other. For each j such that $m+1 \leq j \leq n$, we pick an arbitrary corner (e_j^1, v_j, e_j^2) of X_j. Consider the following two corners of X:

$$((e_1^1, \ldots, e_1^1, e_{m+1}^1, \ldots, e_n^1), (v_1, \ldots, v_1, v_{m+1}, \ldots, v_n), (e_1^2, \ldots, e_1^2, e_{m+1}^2, \ldots, e_n^2)) \text{ and}$$

$$((e_1^{1'}, \ldots, e_1^{1'}, e_{m+1}^1, \ldots, e_n^1), (v_1', \ldots, v_1', v_{m+1}, \ldots, v_n), (e_1^{2'}, \ldots, e_1^{2'}, e_{m+1}^2, \ldots, e_n^2)).$$

By Theorem 7.12, $\mathrm{Aut}(X)$ is generated by automorphisms of prime factors and permutation of isomorphic factors. In particular, it is impossible for $\mathrm{Aut}(X)$ to map one of the above corners to the other, contradicting to the platonicity of X. Therefore we can conclude that any prime factor of X is platonic. □

The corollary below answers the question we posed in the beginning of the chapter.

Corollary 7.14. For $i \in \{1, 2, \ldots, n\}$, let X_i be a simple prime complex with a finite simple connected non-bipartite symmetric R-thin 1-skeleton having more than one vertex. Then the complex tensor product $X = \otimes_{i=1}^n X_i$ has automorphism group $\mathrm{Aut}(X)$ generated by $\mathrm{Aut}(X_i)$'s and permutations of isomorphic X_j's.

Proof. By Proposition 7.9, we know X is a simple complex. By the definition of graph tensor product, we know X^1 is a finite simple graph. Note that a simple graph is non-bipartite if and only if there is a cycle of odd length. Then the graph tensor product of two non-bipartite graphs contains a cycle of odd length and hence is non-bipartite. Induction shows that X^1 is non-bipartite, and by Theorem 3.4 we know that X^1 is connected. By the special case of Theorem 5.15 (platonic complexes without faces), we know X^1 is symmetric and hence edge-transitive. Note that for two graphs Γ_1 and Γ_2, the set of neighbours of a vertex $(u, v) \in V(\Gamma_1 \otimes \Gamma_2)$ is the direct product of the set of neighbours of u in Γ_1 with the set of neighbours of v in Γ_2. This implies the graph tensor product of R-thin graphs is a R-thin graph. To summarize, we know X is a simple complex with a prime factorization $X = \otimes_{i=1}^{n} X_i$, and its 1-skeleton X^1 is a finite simple connected non-bipartite edge-transitive R-thin graph with more than one vertex. By Theorem 7.12, we know that $\mathrm{Aut}(X)$ is as described in the corollary. □

Remark. The tensor products of edge-transitive graphs are not necessarily edge-transitive. Therefore we require each X_i^1 to be symmetric to ensure the edge-transitivity of X^1.

Note that when a complex has a face of odd length, then the 1-skeleton of the complex is non-bipartite, and Corollary 7.14 has a chance to work. In the next chapter, we will investigate the automorphism group of the tensor product of complexes with only faces of even lengths from a different aspect.

Chapter 8

Tensor Symmetry II

In this chapter we investigate the tensor product of complexes with only faces of even lengths, and our goal is to develop results similar to Corollary 7.14, which basically says an automorphism of certain complex tensor products must be of Cartesian type. Note that when there is more than one bipartite factor, Theorem 3.4 implies that the complex tensor product is disconnected, and the product is likely to have non-Cartesian automorphisms from the direct product of automorphism groups of components. Hence in such a context, the proper question to pose should be as follows: for complexes X_i with only faces of even lengths, is the automorphism group of a component of $\otimes X_i$ generated by automorphisms of X_i's together with permutations of isomorphic factors?

For graph tensor products, the example of K_2 × dodecahedron graph in Figure 3.9 shows that the connectedness of the product does not guarantee the absence of non-Cartesian automorphisms. For complex tensor products, we hope that the extra face structure helps to eliminate non-Cartesian automorphisms. For example, let us look at the complex tensor product of two squares, which has two isomorphic components. We denote vertices of a square by $0, 1, 2, -1$ cyclically, and illustrate one component of the product in Figure 8.1. Note that the 1-skeleton of the component is actually a complete bipartite graph with $2 \cdot 4! \cdot 4!$ automorphisms, and not all of them give a complex automorphism due to the extra face structure.

Figure 8.1 also reveals an important fact of the tensor product of complexes with only faces of even lengths: a face is antipodally attached to another face generated by the same pair of faces, and through such antipodally attached relation we can find all other faces generated by the same pair of faces in that component. Such face blocks (defined in Definition 8.4) help to determine the Cartesian structure of a complex tensor product,

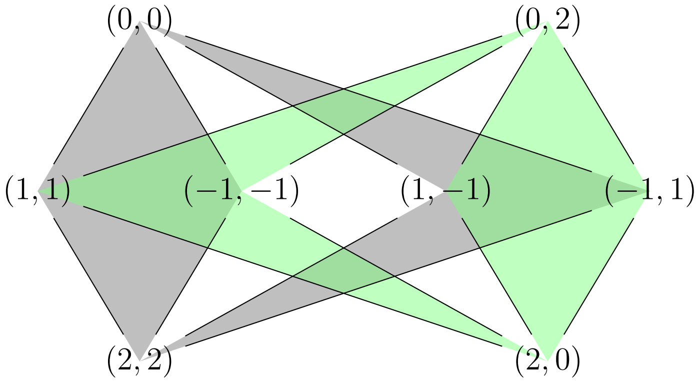

Figure 8.1: a component of the tensor product of two squares

and if we can show a generic face block has only Cartesian automorphisms, then we have a chance to force a complex automorphism stabilizing a face block to be of Cartesian type. To simplify the problem, we restrict our discussion to the tensor product of complexes with faces of the same even length, and the first step is to establish the Cartesian result for the tensor product of $2n$-gons. The following lemma is a useful tool for this purpose.

Lemma 8.1. Suppose on a real line, someone wants to take d steps to walk from an integer $d-2k$ to 0, where $\lfloor \frac{d}{2} \rfloor \geq k \geq 0$ is an integer, and each step is either plus 1 or minus 1. Then there are $\binom{d-1}{k}$ ways to arrive from 1, and $\binom{d-1}{k-1}$ ways to arrive from -1. The ratio $\binom{d-1}{k}/\binom{d-1}{k-1}$ is greater than or equal to 1, with equality if and only if $d-2k=0$. Moreover, when d is fixed and k is increasing, the ratio is decreasing.

Proof. Suppose this person takes x steps of minus 1 and y steps of plus 1 to arrive at 0. Then we have $x+y=d$ and $-x+y=-d+2k$, and therefore $x=d-k$ and $y=k$. By ordering two types of steps arbitrarily, we can obtain all different ways to arrive at 0. To arrive from 1, the last step must be minus 1, and there are $\binom{d-1}{k}$ such combinations. To arrive from -1, the last step must be plus 1, and there are $\binom{d-1}{k-1}$ such combinations. When d is odd, we have $\frac{d-1}{2} \geq k$ and hence $\binom{d-1}{k} > \binom{d-1}{k-1}$. When d is even, we have $\frac{d}{2} \geq k$ which implies $\frac{d-1}{2} > k-1$ and hence $\binom{d-1}{k} \geq \binom{d-1}{k-1}$, with equality if and only if $k+(k-1)=d-1$, namely $d-2k=0$. To show that the ratio $\binom{d-1}{k}/\binom{d-1}{k-1}$ decreases as k increases, we simply have to verify the following inequality:

$$\binom{d-1}{k}/\binom{d-1}{k-1} > \binom{d-1}{k+1}/\binom{d-1}{k}$$

$$\Leftrightarrow \quad \binom{d-1}{k}\binom{d-1}{k} > \binom{d-1}{k+1}\binom{d-1}{k-1}$$

$$\Leftrightarrow \quad \frac{(d-1)\ldots(d-k)}{k!}\frac{(d-1)\ldots(d-k)}{k!} > \frac{(d-1)\ldots(d-k-1)}{(k+1)!}\frac{(d-1)\ldots(d-k+1)}{(k-1)!}$$

$$\Leftrightarrow \quad \frac{d-k}{k} > \frac{d-k-1}{k+1}$$

$$\Leftrightarrow \quad \frac{d}{k} - 1 > \frac{d}{k+1} - 1$$

$$\Leftrightarrow \quad k+1 > k,$$

which is obviously true. □

Proposition 8.2. For $i \in \{1, 2, \ldots, m\}$, let C_i be a graph which is a cycle of length $2n$, where n is an integer at least 3. Then the automorphism group of a component of $\otimes_{i=1}^{m} C_i$ can be generated by elements of $\mathrm{Aut}(C_i)$'s together with permutations of C_i's.

Proof. We denote vertices of C_i by $0, 1, \ldots, n-1, n, -(n-1), -(n-2), \ldots, -1$ cyclically, and let Γ be the component of $\otimes_{i=1}^{m} C_i$ containing the vertex $v = (0, 0, \ldots, 0)$. Note that $\times_{i=1}^{m} \mathrm{Aut}(C_i)$ acts transitively on vertices of $\otimes_{i=1}^{m} C_i$. Therefore to prove this proposition, it suffices to show that the v-stabilizer G_v of $\mathrm{Aut}(\Gamma)$ can be generated by elements of $\mathrm{Aut}(C_i)$'s together with permutations of C_i's. Notice that there are $2^m \cdot m!$ Cartesian automorphisms of Γ fixing v, generated by the reflection fixing 0 in each C_i and all permutations of m factors. If we can show $|G_v| \leq 2^m \cdot m!$, then the proposition follows.

First we show that Γ is a rigid graph. Namely we want to show that if $\varphi \in G_v$ fixes all neighbours of v, then φ must be trivial. Note that two vertices $(b_1, b_2, \ldots, b_m)$ and $(c_1, c_2, \ldots, c_m)$ are adjacent if and only if $b_i - c_i \equiv \pm 1 \mod 2n$ for all i, and therefore

$$V(\Gamma) \subseteq V^* = \{(a_1, a_2, \ldots, a_m) \in V(\otimes_{i=1}^{m} C_i) \mid a_1 \equiv a_2 \equiv \cdots \equiv a_m \mod 2\}.$$

For each $u = (a_1, a_2, \ldots, a_m) \in V^*$, there is a path of length $d = \max\{|a_1|, |a_2|, \ldots, |a_m|\}$ from u to v, because we can reach 0 in d steps in the coordinates with absolute value d, and we can also reach 0 in d steps in the other coordinates by walking back and forth as each coordinate has the same parity. Hence $V(\Gamma) = V^*$, and $d(u, v) = d$ follows easily.

Note that the number of geodesics from u to v is the product of the number of ways in each coordinate to walk to 0 in d steps. Look at the i-th coordinate of v. For now we assume that $a_i \geq 0$, and let k_i be the integer such that $a_i = d - 2k_i$. If $n > a_i > 0$, we have $\lfloor \frac{d}{2} \rfloor \geq k_i \geq 0$, and walking to 0 in d steps is equivalent to the setting of Lemma 8.1. By the lemma, the ratio of numbers of $u - v$ geodesics arriving from 1 and from -1 in the i-th coordinate is $\binom{d-1}{k_i} / \binom{d-1}{k_i - 1} > 1$. Since the automorphism φ fixes $(\pm 1, \pm 1, \ldots, \pm 1)$ and preserves geodesics, this ratio does not change under φ. Again by the Lemma, k_i

must remain the same to keep this ratio, and hence the i-th coordinate of $\varphi(u)$ must be a_i. If $a_i = 0$, then u has a neighbour w with the i-th coordinate 1. Note that $\varphi(u)$ is adjacent to $\varphi(w)$ with the i-th coordinate 1, and the i-th coordinate of $\varphi(u)$ is either 0 or 2. In the latter case, since $n > 2 > 0$, by taking φ^{-1} the above argument implies $a_i = 2$, a contradiction. Hence the i-th coordinate of $\varphi(u)$ is 0. Similarly if $a_i = n$, then the i-th coordinate of $\varphi(u)$ is n. For negative a_i, by applying the mirror version of Lemma 8.1, we know that the i-th coordinate of $\varphi(u)$ is a_i. Note that the above result is true for every coordinate. Hence $\varphi(u) = u$ for every $u \in V(\Gamma)$, and φ is trivial.

Now look at the local structure around v. Note that two neighbours of v taking different values in k coordinates have 2^{m-k} common neighbours. In particular, two neighbours of v differ in exactly one coordinate if and only of they have 2^{m-1} common neighbours. Hence among the neighbours of v, the relation of differing in exactly one coordinate is preserved under G_v. If we draw an auxiliary edge between any two such neighbours of v, then the 2^m neighbours of v plus these auxiliary edges form a hypercube Q_m preserved under G_v. Since Γ is rigid, an automorphism of G_v is completely determined by its action on the neighbours of v, which also induces an automorphism of the auxiliary Q_m. As a result, we have $|G_v| \leq \mathrm{Aut}(Q_m) = 2! \cdot m!$, which finishes the proof. □

Remark. Let H be the subgroup of $\times_{i=1}^{m} \mathbb{Z}_{2n}$ generated by $S = \{(\pm 1, \pm 1, \ldots, \pm 1)\}$. Note that the component Γ in the above proof is actually isomorphic to the Cayley graph of H with respect to the generating set S.

Corollary 8.3. Suppose that X_i is a $2n$-gon for $i \in \{1, 2, \ldots, m\}$, where n is an integer at least 3. Then the automorphism group of a component of $\otimes_{i=1}^{m} X_i$ can be generated by elements of $\mathrm{Aut}(X_i)$'s together with permutations of X_i's.

Proof. Note that a $2n$-gon has the same automorphism group as its 1-skeleton, and $\otimes_{i=1}^{m} X_i$ has the same Cartesian automorphisms as $\otimes_{i=1}^{m} X_i^1$. Hence a vertex stabilizer G_v of a component X of $\otimes_{i=1}^{m} X_i$ has $2^m \cdot m!$ Cartesian automorphisms, and $|G_v|$ is at most the cardinality of the stabilizer of v in X^1, which is $2^m \cdot m!$ by Proposition 8.2. □

Remark. We do need the condition $n \geq 3$ in Proposition 8.2 and Corollary 8.3. For $n = 2$, Figure 8.1 illustrates a component of the tensor product of two squares. Its 1-skeleton is the complete bipartite graph $K_{4,4}$ with lots of non-Cartesian automorphisms. With the face structure, there are much fewer complex automorphisms, but swapping $(0, 2)$ and $(2, 0)$ still gives a non-Cartesian complex automorphism.

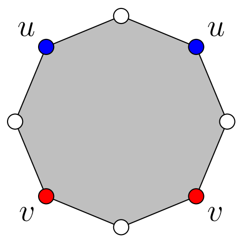

Figure 8.2: a non-elementary complex

Now we formally define the face blocks mentioned in the beginning of the chapter. An intuitive definition of a face block in a complex tensor product $\otimes_{i=1}^{m} X_i$ would be any connected component in $\otimes_{i=1}^{m} f_i$, where each f_i is a face of X_i. Note that if each f_i is an even gon attached injectively, then $\otimes_{i=1}^{m} f_i$ has 2^{m-1} components, and hence 2^{m-1} face blocks. If these f_i's are attached non-injectively, then the above face blocks could have extra incidence relations, and we might end up having fewer components. We would like to define a face block regardless of attaching maps, so we take the following definition.

Definition 8.4. For $i \in \{1, 2, \ldots, m\}$, let X_i be a polygonal cell complex with only faces of even length $2n \geq 2$. Let f_i be a face of X_i with corners labeled by $0, 1, \ldots, 2n-1$ cyclically. A **face block** generated by $f_1, f_2, \ldots, f_m$ is a subcollection of faces generated by $f_1, f_2, \ldots, f_m$ such that two faces f_a and f_b are in the same face block if and only if a corner of f_a with label $(a_1, a_2, \ldots, a_m)$ and a corner f_b with label $(b_1, b_2, \ldots, b_m)$ have

$$a_1 - b_1 \equiv a_2 - b_2 \equiv \cdots \equiv a_m - b_m \mod 2.$$

Remark. It is easy to see that a face block is well-defined no matter how faces are cyclically labeled and no matter which corners are chosen to verify the above criterion. In general it is not obvious whether or not two faces are in the same face block of a complex tensor product without knowing the tensor product structure. In the tensor product of the following class of complexes, recognizing a face block is much easier.

Definition 8.5. A connected polygonal cell complex X is an **elementary** complex if X satisfies the following three conditions:

(1) Every face of X is of the same even length ≥ 2.

(2) No antipodal corners of a face are attached to the same vertex.

(3) For any two vertices, there is at most one pair of antipodal face corners attached.

Remark. Condition (3) basically says no two faces can be attached antipodally, and in a face different pairs of antipodal corners are not attached to the same pair of vertices. For example, the complex in Figure 8.2 is not an elementary complex.

Proposition 8.6. For $i \in \{1, 2, \ldots, m\}$, let X_i be an elementary complex with faces of even length $2n \geq 2$. Then in the complex tensor product $\otimes_{i=1}^m X_i$, for any antipodal vertices u and v of a face in $\otimes_{i=1}^m X_i$, there are exactly 2^{m-1} faces having u and v as antipodal vertices, and these faces are in the same face block. Moreover, for any two faces f and f' in the same face block, we can find a series of faces $f_0, f_1, \ldots, f_k$ such that $f_0 = f$, $f_k = f'$, f_i and f_{i+1} share antipodal vertices for $i \in \{0, 1, \ldots, k-1\}$, and $k \leq n$.

Proof. In $\otimes_{i=1}^m X_i$, suppose that $u = (u_1, u_2, \ldots, u_m)$ and $v = (v_1, v_2, \ldots, v_m)$ are antipodal vertices of a face f generated by $f_1, f_2, \ldots, f_m$, where u_i and v_i are vertices of X_i and f_i is a face of X_i for $i \in \{1, 2, \ldots, m\}$. Note that for each $i \in \{1, 2, \ldots, m\}$, projecting f to X_i gives f_i, and f_i has u_i and v_i as antipodal vertices. Since X_i is elementary, u_i and v_i are not the same vertex, and f_i is the only face of X_i having u_i and v_i as antipodal vertices, with a unique pair of antipodal corners attached to u_i and v_i. Hence any face in $\otimes_{i=1}^m X_i$ having u and v as antipodal vertices must be generated by $f_1, f_2, \ldots, f_m$ in such a way that the corresponding corners c_i of the f_i's at u_i are combined together. With the corner c_1 of f_1 fixed, flipping f_i at c_i for $i \in \{1, 2, \ldots, m\}$ gives all 2^{m-1} faces having u and v as antipodal vertices, and these faces are in the same face block.

Now suppose that f and f' are two faces in the same face block B generated by faces with corners labeled by $0, 1, \ldots, 2n-1$ cyclically. Then we can label corners in B according to such a corner labeling, and by following steps of $(\pm 1, \pm 1, \ldots, \pm 1)$, we can start from a vertex v of f to reach any other vertex in B in n steps. In particular, there is a unique vertex in B such that we need n steps to reach it from v. Since f' has more than one vertex, we can start from v to reach a vertex u of f' in $n-1$ steps. By adding one step in f and one step in f' if necessary, we can find a path from f to f' of length at most $n+1$ such that the first and the last steps are in f and f' respectively. Note that each $(\pm 1, \pm 1, \ldots, \pm 1)$ step determines a unique face in B, and hence the above path determines a series of faces $f_0, f_1, \ldots, f_k$ such that $f_0 = f$, $f_k = f'$, and $k \leq n$. If f_i and f_{i+1} are determined by the same $(\pm 1, \pm 1, \ldots, \pm 1)$ step, then f_i and f_{i+1} are actually the same face, and we can remove one of them from the sequence. If f_i and f_{i+1} are determined by different $(\pm 1, \pm 1, \ldots, \pm 1)$ steps, then f_i and f_{i+1} are two different faces with a common vertex with label $(a_1, a_2, \ldots, a_m)$. Note that

$$(a_1, a_2, \ldots, a_m) + n(\pm 1, \pm 1, \ldots, \pm 1) = (a_1 + n, a_2 + n, \ldots, a_m + n) \mod 2n,$$

which is also a common vertex of f_i and f_{i+1}, and therefore f_i and f_{i+1} share antipodal vertices. The above argument is illustrated in Figure 8.1. $\square$

Proposition 8.6 allows us to easily recognize a face block in a complex tensor product. If we impose the following conditions on each factor, then we can read the Cartesian structure of a complex tensor product through the incidence relation of face blocks.

Definition 8.7. A connected polygonal cell complex X is an **ordinary** complex if every face f of X is of the same even length $2n \geq 4$, and satisfies the following extra conditions:

(1) If we label corners of f cyclically from 1 to $2n$, then any two corners with different parities are not attached to the same vertex.

(2) For any face f' incident to f, either f has only one corner meeting f', or f has only two consecutive corners meeting f'.

Remark. If the 1-skeleton of X is bipartite, then X satisfies (1) automatically. Also note that a polygonal complex satisfies both (1) and (2). The reader might have noticed that (2) implies the condition (3) of an elementary complex. Since there are alternative conditions serving our purpose as effectively as (2), we avoid defining ordinary complexes as a subclass of elementary complexes.

Proposition 8.8. For $i \in \{1, 2, \ldots, m\}$, suppose that X_i is an ordinary complex with faces of even length $2n \geq 4$. Let B be a face block generated by $f_1, f_2, \ldots, f_m$ and B' be a face block generated by $f'_1, f'_2, \ldots, f'_m$, where f_i and f'_i are faces of X_i. If B and B' are incident, then the following two statements are equivalent:

(1) $\exists j$ such that f_j is incident to f'_j in X_j, and $\forall i \neq j$ we have $f_i = f'_i$.

(2) Every face of B is incident to a face of B'.

Proof. Assume (1). Without loss of generality, we can assume that $j = 1$. Since B and B' are incident, there is a face corner c of B meeting a face corner c' of B'. Suppose that c is the combination of corners c_i of the f_i's, and c' is the combination of corners c'_i of the f'_i's. Note that c_1 of f_1 meets c'_1 of f'_1 in X_1. Also note that for $i \neq 1$, c_i and c'_i are in the same face f, and they are either the same corner or different corners attached to

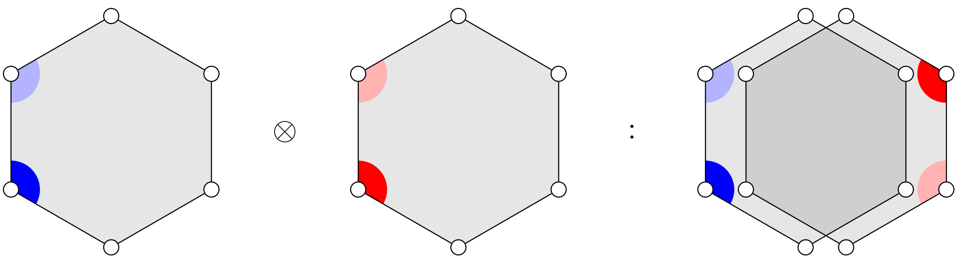

Figure 8.3: how to avoid incidence

the same vertex. In particular, by condition (1) of Definition 8.7, c_i and c_i' have the same parity under cyclic $\mathbb{Z}_{2n}$ labeling for $i \neq 1$. Let f be an arbitrary face of B generated by combining c_1 of f_1 with corners $\overline{c_i}$ of the f_i's for $i \neq 1$. By Definition 8.4, $\overline{c_i}$ has the same parity as c_i, and therefore has the same parity as c_i'. Then again by Definition 8.4, the face f' generated by combining c_1' of f_1' with $\overline{c_i}$'s of the f_i's is a face of B'. It is obvious that f is incident to f'. To summarize, given an arbitrary face f of B, we can find a face f' of B' incident to f. Hence (1) implies (2).

Assume (2). If f_i and f_i' are disjoint, then B and B' are disjoint, which contradicts (2). Hence for each $i \in \{1, 2, \ldots, m\}$, f_i and f_i' are either incident or actually the same. Suppose that there is more than one j, say for $j \in \{1, 2\}$, such that f_j and f_j' are incident. By condition (2) of Definition 8.7, f_1 and f_2 have either one corner or two consecutive corners meeting f_1' and f_2' respectively. Pick two consecutive corners of f_1 containing all corners meeting f_1' and colour them blue. Similarly pick two consecutive corners of f_2 containing all corners meeting f_2' and colour them red. Consider the faces generated by $f_1, f_2, \ldots, f_m$ with the following corner combination: coloured corners of f_1 and f_2 are placed at the opposite positions, as illustrated in Figure 8.3. Note that these faces are disjoint with faces generated by $f_1', f_2', \ldots, f_m'$. If B does not contain any of these faces, we can flip two red corners of f_2 to generate faces of B, and the resulting faces are still disjoint with faces generated by $f_1', f_2', \ldots, f_m'$. In other words, we can find a face of B incident to no face in B', a contradiction. So there is at most one j such that f_j and f_j' are incident. Moreover, condition (1) of Definition 8.7 implies that different face blocks generated by $f_1, f_2, \ldots, f_m$ are disjoint. Since B and B' are incident, we know that there is exactly one j such that f_j and f_j' are incident. Hence (2) implies (1). $\square$

Remark. Note that condition (2) of Definition 8.7 is only used for the argument illustrated in Figure 8.3. It is not hard to have alternative conditions serving this purpose, especially when the length of faces is higher. We also want to point out that through

$$
\begin{tikzcd}
f_1^1 \arrow[d, dash] & B^{1,1} \arrow[r, dash] \arrow[d, dash] & B^{1,2} \arrow[r, dash] \arrow[d, dash] & B^{1,3} \arrow[d, dash] \\
f_1^2 & B^{2,1} \arrow[r, dash] & B^{2,2} \arrow[r, dash] & B^{2,3} \\
 & f_2^1 \arrow[r, dash] & f_2^2 \arrow[r, dash] & f_2^3
\end{tikzcd}
$$

Figure 8.4: Cartesian structure of face blocks

finer examination of incidence relation between face blocks, it is possible to obtain more information such as how f_j meets f_j' in X_j, perhaps under weaker conditions.

With Propositions 8.6 and 8.8, in a tensor product $X = X_1 \otimes X_2 \otimes \cdots \otimes X_m$ where each X_i is an elementary ordinary complex with only faces of even length $2n \geq 4$, we can recognize face blocks and the Cartesian structure of X through the incidence relation on faces, which is preserved under automorphisms of X. Now we define a graph Γ_X to encode the Cartesian structure of X. Let Γ_X be a simple graph with vertex set $\times_{i=1}^{m} F(X_i)$, where a vertex $(f_1, f_2, \ldots, f_m)$ represents all faces of X generated by $f_1, f_2, \ldots, f_m$, such that two vertices are adjacent if and only if they take the same face in $m-1$ coordinates, and have incident faces in the remaining coordinate. Let Γ_{X_i} be a simple graph with vertex set $F(X_i)$, such that two vertices are adjacent if and only if the corresponding faces are incident in X_i. Notice that $\Gamma_X = \Gamma_{X_1} \square \Gamma_{X_2} \square \cdots \square \Gamma_{X_i}$. Figure 8.4 illustrates the case $m = 2$, where $B^{i,j} = (f_1^i, f_2^j)$ represents all faces generated by f_1^i and f_2^j. The following theorem due to Imrich [19] and Miller [25] restricts the automorphism group of Γ_X.

Theorem 8.9. Suppose that Γ is a finite simple connected graph with a factorization $\Gamma = \Gamma_1 \square \Gamma_2 \square \cdots \square \Gamma_m$, where each Γ_i is prime with respect to Cartesian product. Then the automorphism group of Γ is generated by automorphisms of prime factors and permutations of isomorphic factors.

We can not guarantee Γ_{X_i} is prime, but at least X_i is indeed a prime complex.

Proposition 8.10. Let Y be an elementary complex. Then Y is a prime with respect to complex tensor product, and Y is not a component of any complex tensor product.

Proof. Suppose that there exist complexes Y_1 and Y_2 such that Y is a component of $Y_1 \otimes Y_2$. Note that a face of Y is of even length, and must be generated by either two

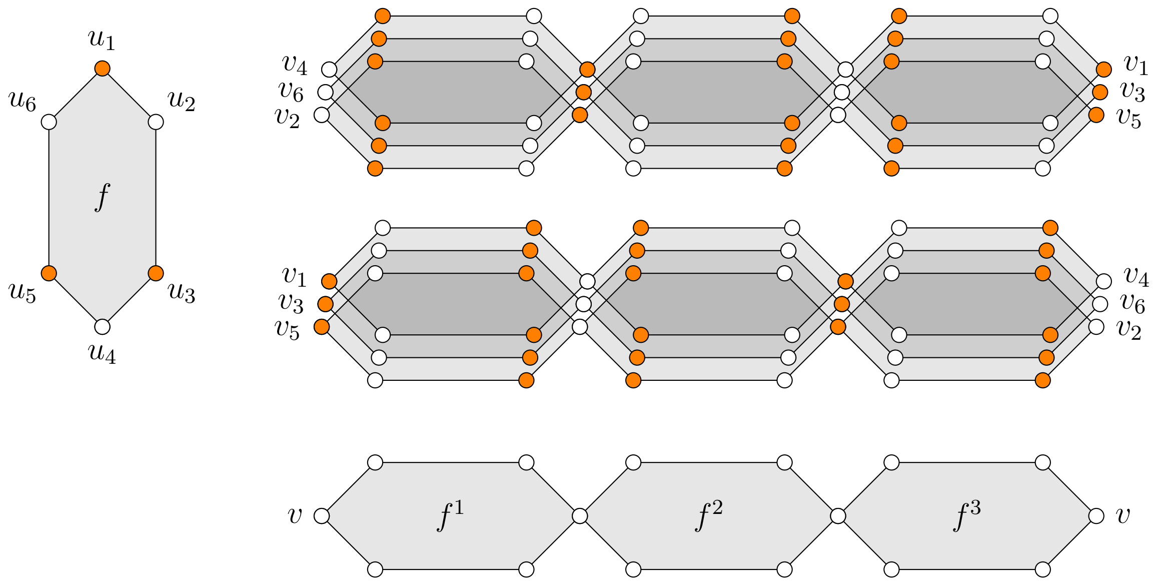

Figure 8.5: tensor product of a hexagon with a 3-hexagon necklace

even faces or by one even and one odd face. In either case, by Definition 5.7, Y will have faces antipodally attached together, violating that Y is elementary. $\square$

Note that in Figure 8.4, each $B^{i,j}$ actually contains two face blocks generated by f_1^i and f_2^j, and in general each vertex of Γ_X defined above contains 2^{m-1} face blocks. Even if we have some control over the automorphism group of Γ_X, having multiple face blocks at one vertex of Γ_X could lead to non-Cartesian automorphisms of X. Let us look at the tensor product of a hexagon with a 3-hexagon necklace as illustrated in Figure 8.5, where v_i is the vertex generated by u_i and v, and coloured vertices in the product are generated by coloured u_1, u_3, and u_5. For brevity, half of the faces in the product are omitted. Consider the automorphism ρ of the product induced by fixing f, f^1, and f^3 but flipping f^2 (swapping the top and the bottom edges) in two factors. Then ρ fixes the four face blocks on the left and right, and permutes vertices in each of the two middle blocks. In particular, we can permute vertices in a block while its two incident blocks are fixed. Therefore we can permute vertices in one middle block and fix all other five blocks. This gives a non-Cartesian automorphism.

There are two main reasons why we have the above non-Cartesian automorphism. First, there is more than one face block generated by the same faces lying in the same component of the product. Secondly, factors are not rigid enough, so the action on one face block can not affect incident blocks, and can not be transmitted to blocks generated by the same faces. We suspect that if either of these two reasons is absent, then each component of the product might have only Cartesian automorphisms. In particular, if the 1-skeleton of each factor is bipartite, then face blocks generated by the same faces are in different components. Also note that if a complex is a surface, it is rigid enough that the action on one face completely determines the whole automorphism. So far we do not have a definite result yet, and hence we pose the following two conjectures. We hope to resolve these problems in the near future.

Conjecture 8.11. For $i \in \{1, 2, \ldots, m\}$, suppose that X_i is an elementary ordinary complex with faces of the same even length $2n \geq 6$, and X_i has bipartite 1-skeleton. Then for any component X of the complex tensor product $\otimes_{i=1}^{m} X_i$, $\mathrm{Aut}(X)$ can be generated by automorphisms of X_i's together with permutations of isomorphic factors.

Conjecture 8.12. For $i \in \{1, 2, \ldots, m\}$, suppose that X_i is an elementary ordinary complex with faces of the same even length $2n \geq 6$, and X_i has surface structure. Then for any component X of the complex tensor product $\otimes_{i=1}^{m} X_i$, $\mathrm{Aut}(X)$ can be generated by automorphisms of X_i's together with permutations of isomorphic factors.

Chapter 9

Almost Platonic Complexes

First we introduce some terminology from graph theory. An **s-arc** in a graph is a sequence $(v_0, e_1, v_1, e_2, \dots, e_s, v_s)$ of vertices v_i and edges e_j such that e_i connects v_{i-1} and v_i for each $i \in \{1, \dots, s\}$, and $e_j \neq e_{j+1}$ for each $j \in \{1, \dots, s-1\}$. For example, a 1-arc is essentially an edge with an orientation. We can also think of a 1-arc as an edge-vertex flag in a graph. A graph is **s-transitive** if its automorphism group acts transitively on s-arcs, but not on $(s+1)$-arcs. When there is no confusion, especially for graphs without parallel edges, we denote an s-arc simply by its vertices $(v_0, v_1, \dots, v_s)$.

A graph is said to be **half-transitive** if its automorphism group acts transitively on vertices and edges, but not on 1-arcs. Using the language of flags, the automorphism group of a half-transitive graph acts transitively on each type of partial flag, but not on flags. The smallest half-transitive graph is known as the Holt graph [18]. In this chapter we discuss the analogous property for complexes.

Definition 9.1. A polygonal cell complex is **almost platonic** if its automorphism group acts transitively on each type of partial flag, but not on fev flags.

Proposition 9.2. Suppose X is a polygonal cell complex. X is almost platonic if and only if the following three conditions hold:

(1) $\mathrm{Aut}(X)$ acts transitively on ev and f flags.

(2) For each face in X, there exists $\alpha \in \mathrm{Aut}(X)$ rotating the face by 1.

(3) If $\beta \in \mathrm{Aut}(X)$ swaps the endpoints of an edge, then β fixes no incident faces.

Proof. Suppose X satisfies (1), (2), and (3). By rotation of faces and the transitivity on faces, we get the transitivity on fe and fv flags. (3) prevents $\mathrm{Aut}(X)$ from acting transitively on fev flags. Therefore X is almost platonic.

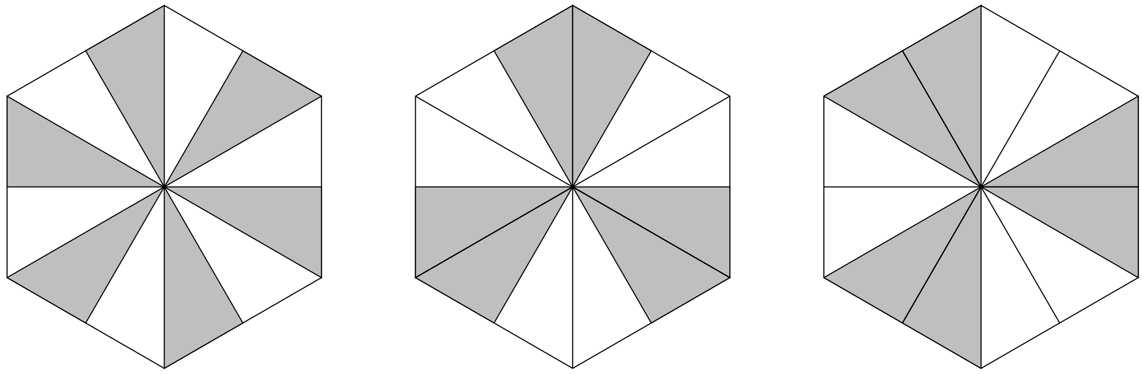

Figure 9.1: three types of partial transitivity

Conversely, suppose X is almost platonic. (1) is immediate from the definition. Note that every fe flag lies in two fev flags. By the transitivity on fe flags and the non-transitivity on fev flags, $\mathrm{Aut}(X)$ has exactly two orbits of fev flags, and two fev flags containing the same fe flag must be in different orbits. Similarly, two fev flags containing the same fv flag must be in different orbits. The two orbits of fev flags are as illustrated on the left of Figure 9.1. The configuration of fev orbits assures (2) and (3). □

Remark. With the f-transitivity in (1), we can modify (2) as "for some face in X". More importantly, (3) implies that the action of an edge stabilizer on incident faces has 2 blocks of imprimitivity. In particular the edge valency of X must be even.

Proposition 9.3. Suppose X is a polygonal cell complex, and $\mathrm{Aut}(X)$ acts transitively on fe or fv flags, but not on fev flags. Then X is exclusively of one of the following three types as illustrated in Figure 9.1:

(1) $\mathrm{Aut}(X)$ can rotate a face by 1.

(2) $\mathrm{Aut}(X)$ can rotate a face by 2, and reflect a face at a corner.

(3) $\mathrm{Aut}(X)$ can rotate a face by 2, and reflect a face at an edge.

Moreover, if X is ev-transitive and has odd face length, then X is almost platonic.

Proof. If $\mathrm{Aut}(X)$ acts transitively on fe and fv flags, the same argument as in the proof of Proposition 9.2 shows that $\mathrm{Aut}(X)$ can rotate a face by 1. Suppose X is fe-transitive but not fv-transitive. $\mathrm{Aut}(X)$ has exactly two orbits of fev flags, and two fev flags containing the same fe flag must be in different orbits. Non-transitivity on fv flags rules out the existence of rotation by 1, and therefore two fev flags containing the same fv flag must be in the same orbit. The two orbits of fev flags are as illustrated in the middle of Figure 9.1, which shows X is of the type (2). If X is fv-transitive but not fe-transitive, the same reasoning shows that X is of type (3), as the right of Figure 9.1.

Note that faces must have even length in (2) and (3). If X is ev-transitive and has odd face length, then X is certainly of type (1), and therefore $\mathrm{Aut}(X)$ is transitive on fe and fv flags. Since X is not fev-transitive, we know X is almost platonic. □

Now we try to find examples of almost platonic complexes. Suppose the complex X is almost platonic, and for every 2-arc in X, there is at most one face corner attached. Let X^1 be the 1-skeleton of X. Since X is ev-transitive, $\mathrm{Aut}(X)$ acts transitively on 1-arcs of X^1. If $\mathrm{Aut}(X)$ acts transitively on 2-arcs of X^1, then $\mathrm{Aut}(X)$ can swap two edges at a corner of a face in X, and therefore acts transitively on fev flags of X, which is impossible. Note that X^1 as a graph might have more automorphisms than $\mathrm{Aut}(X)$, and therefore X^1 is not necessarily a 1-transitive graph. But for the purpose of finding examples, we can start from assuming that X^1 is 1-transitive. For the case of valency 3, we have the following useful proposition.

Proposition 9.4. Suppose Γ is a connected 1-transitive graph of valency 3. Then for any two 1-arcs in Γ, there is a unique automorphism of Γ sending one arc to the other.

Proof. The existence is from the definition of 1-transitivity. To show the uniqueness, it suffices to show that an automorphism γ of Γ fixing a 1-arc (u, v) must be trivial. Note that γ must fix the other two neighbours of v, similarly for u. Otherwise, the automorphism group will act transitively on 2-arcs, which is impossible. By induction and the connectedness of Γ, we know that γ fixes every vertex in Γ. □

The first known example of 1-transitive graph is given by Frucht [13]. It is constructed as a Cayley graph as follows. Let a_1 be the permutation $(1, 2)(3, 4)(5, 6)$ and $b = (1, 2, 3)(4, 5, 7)(6, 8, 9)$. Let $a_2 = ba_1b^{-1}$ and $a_3 = ba_2b^{-1}$. Note that $a_1 = ba_3b^{-1}$ since b has order 3. Let G be the group generated by $\{a_1, a_2, a_3\}$. It is a group of order 432. The key property for this generating set is that the group automorphism acts transitively on a_1, a_2, a_3, but there is no automorphism fixing a_1 while swapping a_2 and a_3.

Let Γ_1 be the Cayley graph of G with respect to the generating set $\{a_1, a_2, a_3\}$. Any two vertices v and v' are adjacent if and only if $v' = v * a_i$ for some i, and we label this edge with a_i. Note that each a_i has order 2, so we can think of this Cayley graph as an undirected graph of valency 3. With the Cayley graph structure, this graph comes

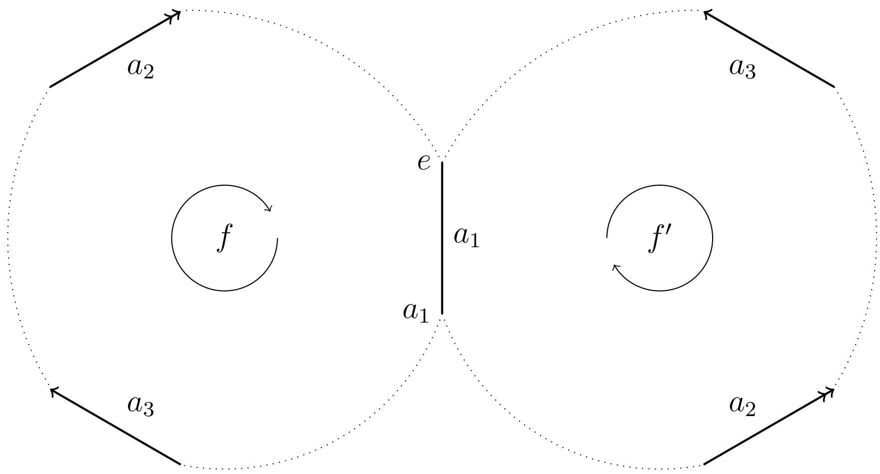

Figure 9.2: neighbourhood of e in X_1

with 2 types of automorphisms. One is by left multiplying each vertex with an element of G, which preserves the labels of edges. Another type of automorphisms is by taking conjugation by some power of b, which fixes the identity e and rotates a_1, a_2, a_3. In [13], Frucht shows that this graph is 1-transitive. By Proposition 9.4, for any two 1-arcs in the graph, there is a unique automorphism sending one arc to the other. By looking at the image of the (e, a_1) arc, each graph automorphism can be uniquely expressed as a conjugation by some power of b followed by a left multiplication. Therefore the automorphism group $\mathrm{Aut}(\Gamma_1)$ can be expressed as a semidirect product $G \rtimes C_3$, which has order $432 \cdot 3$.

Now we try to attach faces to Γ_1. Starting from a vertex $g \in G$, follow the edge with label a_1,a_2, a_3, a_1, a_2, a_3, a_1,.... After 24 steps, we will go back to g for the first time, and this gives a simple closed 24-arc. We attach a face along each simple closed 24-arc of this type to construct a complex X_1. Note that each element in $\mathrm{Aut}(\Gamma_1)$ preserves such 24-arcs, so the automorphism group of X_1 is exactly $\mathrm{Aut}(\Gamma_1)$. It is easy to see that each 2-arc determines a unique face, each vertex lies in 3 faces, and each edge lies in two faces. Therefore X_1 is actually a surface.

Consider the edge $\{e, a_1\}$ and the two incident faces f and f' as shown in Figure 9.2. All 6 dotted arcs have length 7. Program 3 in the appendix verifies that f and f' meet at 3 edges with labels a_1, a_2, and a_3 respectively. So this is a polygonal cell complex, not a polygonal one. The order $a_1, a_2, a_3, a_1, a_2, a_3, \ldots$ gives an orientation to each face of X_1. Note that the orientations of f and f' are compatible. For every edge with label a_i in X_1, there is a left-multiplication sending this edge to the edge in the figure with label a_i while preserving the orientations of two incident faces. This shows that X_1 is an orientable surface. Standard counting argument shows that X_1 has 432 vertices, 648 edges, 54 faces, and therefore genus 82. Moreover, this surface has the property we want.

Proposition 9.5. The polygonal cell complex X_1 is almost platonic.

Proof. From the definition of X_1, its automorphism group $\mathrm{Aut}(\Gamma_1)=G \rtimes C_3$ acts transitively on ev flags. Since left multiplication preserves edge labels, it is easy to see that $\mathrm{Aut}(\Gamma_1)$ acts transitively on faces. Next, we check if $\mathrm{Aut}(\Gamma_1)$ can rotate the face f by 1 in Figure 9.2. This can be done by conjugation by b followed by left multiplication by a_1. A vertex v in f is of the form $a_1a_2a_3a_1a_2a_3$.... Conjugation by b sends this vertex to $a_2a_3a_1a_2a_3a_1$..., and then left multiplication by a_1 gives $a_1a_2a_3a_1a_2a_3a_1$..., the next vertex of v under clockwise rotation.

Since Γ_1 is a connected 1-transitive graph of valency 3, by Proposition 9.4, there is a unique $\gamma \in \mathrm{Aut}(\Gamma_1)$ swaps e and a_1. We see that γ is exactly the left multiplication by a_1, which swaps e and a_1 but sends f to f'. We have shown that X_1 meets all three conditions in Proposition 9.2, and therefore X_1 is almost platonic. □

Remark. Since X_1 has the structure of a surface, we can get another almost platonic complex made of triangles simply by taking the dual complex of X_1.

Corollary 9.6. Platonicity is strictly stronger than transitivity on all partial flags.

Proof. This follows immediately from the existence of an almost platonic complex. □

Corollary 9.7. The edge valency of an almost platonic complex is even. Conversely, for any even number n, there exists an almost platonic complex with edge valency n.

Proof. The first half is in the remark after Proposition 9.2. Note that by duplicating each face of X_1 with the same multiplicity, the new complex we get is still almost platonic. By controlling the multiplicity of face duplication, we can achieve any even edge valency. □

What if we attach different faces to Γ_1? Starting from a vertex in Γ_1, follow the edges with label $a_i, a_j, a_i, a_j, a_i, ...$ This gives a simple closed 12-arc whenever $i \neq j$. By attaching a face along each such 12-arc, we construct another complex X_1'. Note that any graph automorphism of Γ_1 preserves such 12-arcs, so the automorphism group of X_1' is again $\mathrm{Aut}(\Gamma_1)$. It is easy to see that $\mathrm{Aut}(\Gamma_1)$ acts transitively on f and ev flags.

Consider the face with vertex e and labels a_1 and a_2. By Proposition 9.4, there is a unique $\gamma \in \text{Aut}(\Gamma_1)$ sending the 1-arc (e, a_1) to (a_1, a_1a_2). Conjugation by b followed by left multiplication by a_1 does the job. This is not a rotation of the face, and therefore X_1' has no automorphism of rotation by 1 along a face. However left multiplication by a_1a_2 rotates this face by 2, and left multiplication by a_1 swaps e and a_1 while preserving the face. Hence X_1' is of type (3) of Proposition 9.3. X_1' has 2 orbits for fev flags, 2 orbits for fe flags, and is transitive on any other partial flags. Note that X_1' is also a surface, and the dual complex of X_1' has 2 orbits for fev flags, 2 orbits for ev flags, and is transitive on any other partial flags. This dual complex has triangular faces, and is of type (1).

Now we try to find a smaller 1-transitive graph of valency 3. From the aspect of its automorphism group G, G should be a transitive group on n vertices, where the stabilizer of a vertex v has a suborbit $\{x, y, z\}$ of length 3 corresponding to its neighbours. Moreover, if $g \in G$ fixes v and x, then g is trivial. The orbit of the edge $\{v, x\}$ under the action of G gives all edges of the graph, and therefore we can reconstruct the graph from the group G. The problem is, if we start from a group with the property above, the graph constructed might end up having more automorphisms than G. The alternating group A_4 is an example. The corresponding graph under this construction is the 1-skeleton of a tetrahedron, of which the automorphism group is S_4. If the automorphism group of the derived graph has size exactly $|G|$, then this graph is 1-transitive and of valency 3.

Program 4 in the appendix checks all transitive groups on up to 30 letters, and it turns out that there is only one group with the desired property. It is a transitive group on 26 letters. While one letter is fixed, this group has 4 suborbits of length 3 satisfying the requirement. Program 5 verifies that the graphs derived from these 4 suborbits are isomorphic. This 1-transitive graph is known as the F26 graph [7]. Figure 9.3 is a particular drawing of the graph, with vertex set $\mathbb{Z}_{26}$. Two vertices in $\mathbb{Z}_{26}$ are adjacent if and only if they either differ by 1, or one vertex v is odd and the other is $v + 7$.

It is not clear which closed arcs should be faces in Figure 9.3. We can rearrange the layout to get Figure 9.4, verified in Program 5 as well, and now the choice of faces is obvious. This graph tessellates a torus by 13 hexagons, and we choose these 13 hexagons to be faces. These faces are exactly of the form $(m-1, m, m+1, m-6, m-7, m-8)$ for odd m's, which looks like a bow tie of height 2 in Figure 9.3. We call this graph Γ_2, and the resulting complex X_2. Note that X_2 is a polygonal complex, not just a polygonal cell complex like X_1.

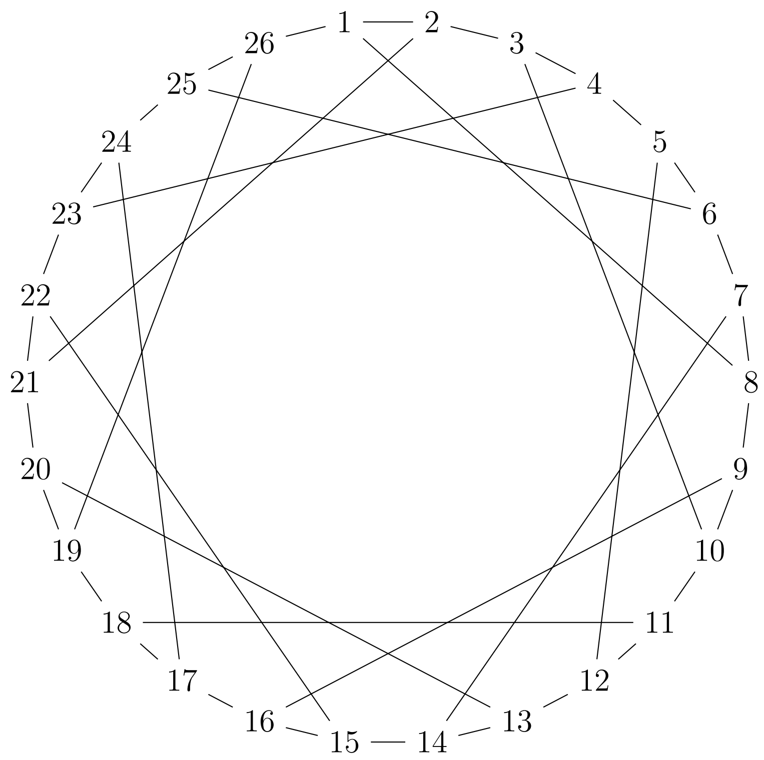

Figure 9.3: a drawing of the F26 graph

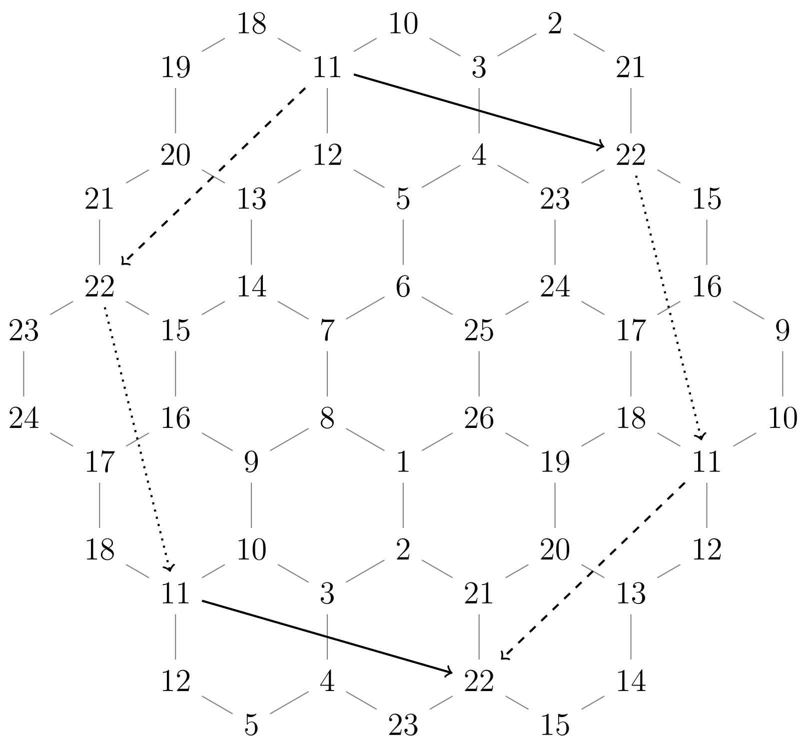

Figure 9.4: another drawing of the F26 graph

Proposition 9.8. The polygonal complex X_2 is almost platonic.

Proof. We look at $\mathrm{Aut}(X_2)$, a subgroup of $\mathrm{Aut}(\Gamma_2)$. In Figure 9.4, we can rotate around the central face while preserving faces. In Figure 9.3, we can rotate around the big circle by 2 while preserving faces. The combination of these two actions shows that $\mathrm{Aut}(X_2)$ is transitive on v, e, f, fe, and fv flags. Can $\mathrm{Aut}(X_2)$ flip an edge? Since Γ_2 is connected, 1-transitive, and of valency 3, by Proposition 9.4, there is a unique element $\gamma \in \mathrm{Aut}(\Gamma_2)$ flipping the edge $\{1, 2\}$, which can be realized as a reflection about the central vertical line in Figure 9.3. Note that γ preserves faces, but swaps two faces containing $\{1, 2\}$. Hence $\mathrm{Aut}(X_2)$ acts transitively on ev flags, but has two orbits for fev flags. $\square$

The above two examples suggest a general method of constructing an almost platonic complex from an arbitrary connected 1-transitive graph Γ of valency 3. Suppose (v_0, v_1, v_2) is an 2-arc in Γ. By Proposition 9.4, there exists a unique automorphism γ sending (v_0, v_1) to (v_1, v_2). The orbit of (v_0, v_1) under the action of γ gives a simple closed n-arc $(v_0, v_1, v_2, \ldots, v_n, v_0)$. Note that any 2-arc (v_i, v_{i+1}, v_{i+2}) determines the same closed arc, and any 2-arc (v_{i+2}, v_{i+1}, v_i), under the action of γ^{n-1}, also determines the same closed arc with the opposite orientation. For each such simple closed arc of Γ (two closed arcs with opposite orientations are identified), we attach a face to Γ along the closed arc, and denote the resulting complex by X_Γ. Note that X_Γ is actually a surface.

Theorem 9.9. Suppose Γ is a connected 1-transitive graph of valency 3. Then the polygonal cell complex X_Γ described above is almost platonic.

Proof. Suppose $(v_0, v_1, \ldots, v_n, v_0)$ is a face of X_Γ, and g is an automorphism of Γ. By the 1-transitivity of Γ and the definition of X_Γ, there exists $\gamma \in \mathrm{Aut}(\Gamma)$ such that $\gamma(v_i) = v_{i+1}$ for each i. Look at the image of $(v_0, v_1, \ldots, v_n, v_0)$ under g. Note that $g\gamma g^{-1} \in \mathrm{Aut}(\Gamma)$ rotates $(g(v_0), g(v_1), \ldots, g(v_n), g(v_0))$. Therefore, $(g(v_0), g(v_1), \ldots, g(v_n), g(v_0))$ is again a face in X_Γ. This shows that every automorphism of Γ preserves faces of X_Γ, and therefore we have $\mathrm{Aut}(X_\Gamma) = \mathrm{Aut}(\Gamma)$.

Now we examine the unique automorphism $\delta \in \mathrm{Aut}(X_\Gamma)$ swapping v_0 and v_1. Suppose the neighbours of v_0 and v_1 other than v_i's are u_0 and u_1 respectively, as shown in Figure 9.5. Assume that δ reflects the face $(v_0, v_1, \ldots, v_n, v_0)$. Then δ swaps v_n and v_2, as well as u_0 and u_1. Note that $\delta * \gamma \in \mathrm{Aut}(X_\Gamma)$ maps (u_0, v_0, v_1) to (u_0, v_0, v_n). This allows $\mathrm{Aut}(X_\Gamma)$ to act transitively on 2-arcs, a contradiction. Therefore δ can not be a reflection of the face. Instead, δ swaps the two faces incident to the edge $\{v_0, v_1\}$. Combined

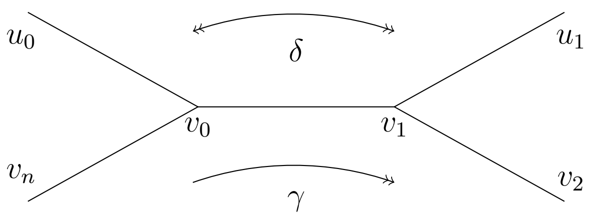

Figure 9.5: when there is a face reflection

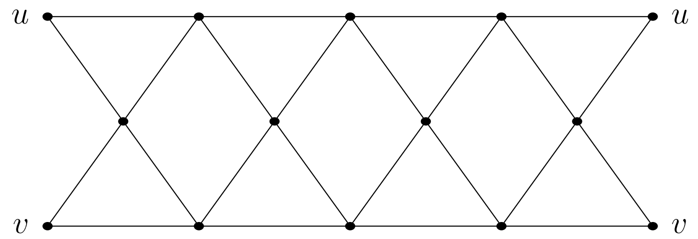

Figure 9.6: 1-skeleton of a cuboctahedron

with ev-transitivity, we know that X is f-transitive. X_Γ satisfies the three conditions of Proposition 9.2. Therefore X_Γ is indeed almost platonic. $\square$

There are two obstacles when we apply this construction to 1-transitive graphs of valency higher than 3. First, the resulting complex might not be f-transitive. This can be circumvented by choosing only one orbit of faces. Secondly, the resulting complex might end up being platonic. For valency 3 cases, the complex can not be platonic mainly because of the lack of corner reflections, a result of 1-transitivity. For higher valency cases, 1-transitivity can not guarantee the construction free of corner reflections. The graph in Figure 9.6 exhibits these two obstacles simultaneously. It can be viewed as the 1-skeleton of a cuboctahedron, and the 1-transitivity follows easily from this visualization. Applying the above construction builds a cuboctahedron with 4 inner hexagonal faces. By only choosing triangular, square, or hexagonal faces, we can obtain an f-transitive complex, but each has corner reflections and therefore is platonic. Note that this graph has a unique automorphism between any two 1-arcs, so such uniqueness is not the key for the construction to work. The real key is the inability of a vertex stabilizer to swap a particular pair of neighbours. Figure 9.6 is actually the worst scenario, where the vertex stabilizer can swap any pair of neighbours.

Suppose Γ is a 1-transitive graph, and there exists a 2-arc (u, v, w) in Γ such that the stabilizer of v can not swap u and w. Because of 1-transitivity, there exists a (not

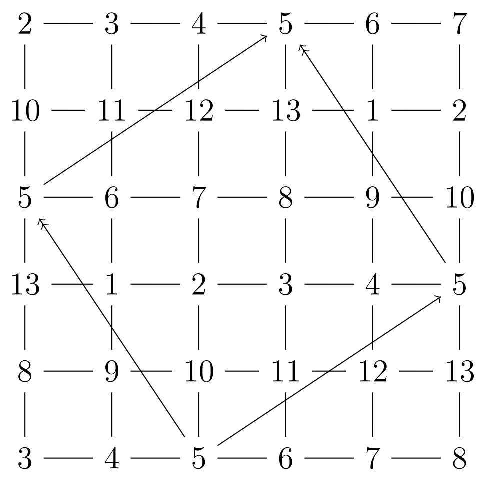

Figure 9.7: general construction for higher valency

necessarily unique) $\gamma \in \mathrm{Aut}(\Gamma)$ sending (u, v) to (v, w). Under the action of Γ, this 2-arc extends to a simple closed n-arc. We attach a face to this n-arc, as well as any image of this n-arc under the action of $\mathrm{Aut}(\Gamma)$. We denote the resulting complex by $X^{\gamma}_{(u,v,w)}$. When such a γ is unique, we can simply write $X_{(u,v,w)}$. In that case, the edge valency of $X_{(u,v,w)}$ is 2, and the neighbourhood of a vertex is homeomorphic to the wedge sum of several open disks, a phenomenon which does occur for almost platonic complexes. For example, duplicating edges and faces but not verticecs of X_1 in Proposition 9.5 gives an almost platonic complex with such topological features, although its 1-skeleton has more than one automorphism between two 1-arcs.

Theorem 9.10. Suppose Γ is a 1-transitive graph, and there exists a 2-arc (u, v, w) in Γ such that the stabilizer of v can not swap u and w. Let $\gamma \in \mathrm{Aut}(\Gamma)$ be an automorphism sending (u, v) to (v, w). Then the complex $X^{\gamma}_{(u,v,w)}$ is almost platonic.

Proof. Since there is at most one face attached to a simple closed arc, the automorphism group of the complex can be viewed as a subgroup of $\mathrm{Aut}(\Gamma)$. Moreover, every element in $\mathrm{Aut}(\Gamma)$ preserves faces of this complex. Therefore the automorphism group of the complex is exactly $\mathrm{Aut}(\Gamma)$. From the construction of the complex, we know it is transitive on *ev* and *f* flags, and $\mathrm{Aut}(\Gamma)$ can rotate a face of the complex by 1. These conditions imply that this complex is transitive on each type of partial flag. Note that $\mathrm{Aut}(\Gamma)$ can not flip the face corner (u, v, w), and therefore this complex can not be platonic. $\square$

Figure 9.7 is an example of the general construction. It is actually a 1-transitive graph such that for any two 1-arcs, there is a unique automorphism sending one arc to the other.

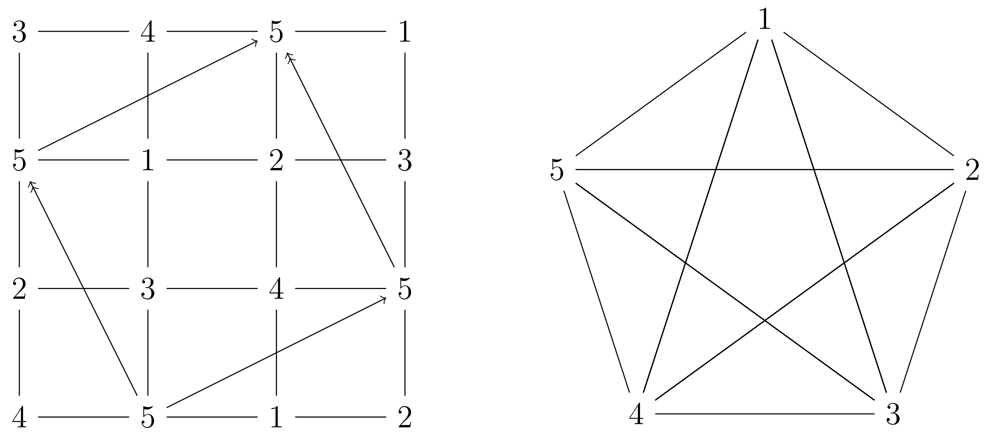

Figure 9.8: almost platonic complex T_1 with 2-transitive 1-skeleton

Note that the graph automorphism group can rotate around a vertex by $\frac{\pi}{2}$, and there is a unique automorphism sending $(2,1)$ to $(2,7)$. Hence the stabilizer of the vertex 2 is exactly C_4, which can not swap neighbours except {1,3} and {7,10}. It is easy to see that $X_{(1,2,7)}$, $X_{(7,2,3)}$, $X_{(3,2,10)}$, $X_{(10,2,1)}$ all give the same tessellation of a torus by 13 squares, and the complex is almost platonic by Theorem 9.10. What if we choose the 2-arc $(1,2,3)$ to construct the complex? The simple closed arc we get is of length 13, and the complex $X_{(1,2,3)}$ has two 13-gons as its faces. With the ability to flip the face corner $(1,2,3)$, $X_{(1,2,3)}$ is platonic. We make this statement explicit in the following proposition.

Proposition 9.11. Suppose X is a platonic complex with a face corner (u,v,w). Then its 1-skeleton X^1 is transitive on 1-arcs, and the stabilizer of v in X^1 can swap u and w. Conversely, suppose Γ is a graph transitive on 1-arcs, and the stabilizer of a vertex v can swap its two neighbours u and w. Let $\gamma \in \mathrm{Aut}(\Gamma)$ be an automorphism sending (u,v) to (v,w). Then the complex $X^{\gamma}_{(u,v,w)}$ is platonic.

Proof. By forgetting the action of $\mathrm{Aut}(X)$ on faces, the quotient of $\mathrm{Aut}(X)$ induces a subgroup of $\mathrm{Aut}(X^1)$, and the first half follows easily. For the second half, the same argument as in Theorem 9.10 shows that the automorphism group of $X^{\gamma}_{(u,v,w)}$ is exactly $\mathrm{Aut}(\Gamma)$, and $X^{\gamma}_{(u,v,w)}$ is transitive on each type of partial flag. Note that there exists $\tau \in \mathrm{Aut}(\Gamma)$ fixing v and swapping u and w. Then τ as a complex automorphism flips the face corner (u,v,w), and this makes $X^{\gamma}_{(u,v,w)}$ platonic. $\square$

Figure 9.4 and Figure 9.7 suggest that the quotient of a plane tessellated by squares or hexagons is a good source of almost platonic complexes. Here we examine two smaller

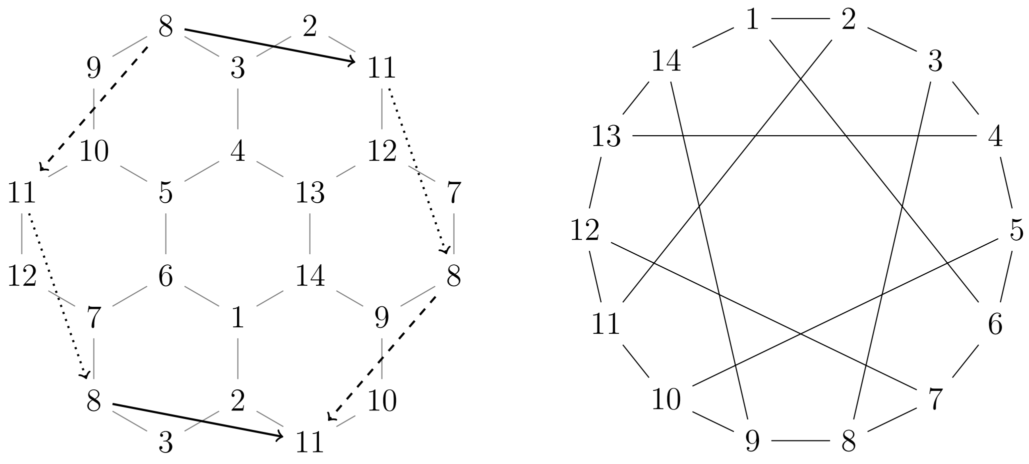

Figure 9.9: almost platonic complex T_2 and the Heawood Graph

examples. Let T_1 be the complex with 5 vertices and 5 square faces as illustrated on the left of Figure 9.8. T_1 is a polygonal cell complex, not a polygonal one. Note that $\mathrm{Aut}(T_1)$ can rotate around a vertex by $\frac{\pi}{2}$, and also rotate around a square by $\frac{\pi}{2}$. If we rotate round 1 by $\frac{\pi}{2}$ and then rotate around the square $(1, 2, 4, 3)$ by $\frac{\pi}{2}$, we obtain a complex automorphism swapping $(1, 3)$. By combining these two types of rotations and edge flipping automorphisms, we know that T_1 is transitive on each type of partial flag. Suppose $\gamma \in \mathrm{Aut}(T_1)$ flips the $(1, 2, 4)$ face corner. Then γ should fix every vertex on the extension of the diagonal through 2 and 3. In turn γ fixes every vertex, a contradiction. Hence there is no corner reflection in $\mathrm{Aut}(T_1)$, and T_1 is almost platonic. The 1-skeleton of T_1 is actually the complete graph of 5 vertices as shown on the right of Figure 9.8. Note that K_5 is 2-transitive. This shows that the 1-skeleton of an almost platonic complex is not necessarily 1-transitive.

Let T_2 be the complex with 14 vertices and 7 hexagonal faces as illustrated on the left of Figure 9.9. As well as T_1 above, T_2 is homeomorphic to a torus, and T_2 is a polygonal complex. Note that $\mathrm{Aut}(T_2)$ can rotate around a vertex by $\frac{2\pi}{3}$, and also rotate around a hexagon by $\frac{\pi}{3}$. By the same argument as T_1, we can show that T_2 is almost platonic. The 1-skeleton of T_2 can be rearranged as the right of Figure 9.9, and the faces are exactly of the form $(m-1, m, m+1, m-4, m-5, m-6)$ for odd m's, which looks like a bow tie of height 2 on the right of the figure. This graph is known as the Heawood graph, which is a 4-transitive graph [16]. Suppose α is a graph automorphism sending $(5, 6, 1, 14, 13)$ to $(13, 14, 1, 6, 5)$, so α flips the closed 6-arc $(4, 5, 6, 1, 14, 13, 4)$. Direction examination shows that α is the permutation $(5, 13)(6, 14)(7, 9)(10, 12)$. So α sends the closed 6-arc

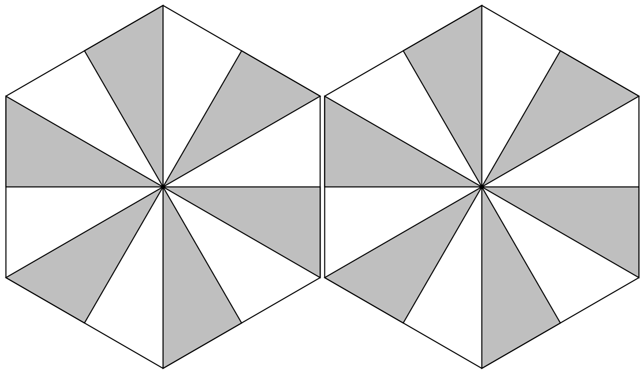

Figure 9.10: orientability of an almost platonic surface

$(1, 2, 3, 8, 7, 6, 1)$, which is a face in T_2, to another closed 6-arc $(1, 2, 3, 8, 9, 14, 1)$, which is no longer a face in T_2. Therefore α is simply a graph automorphism, not a complex automorphism.

So far we have several examples of almost platonic complexes which have surface structure. All these examples are orientable surfaces of genus at least one. The following theorem explains such phenomenon.

Theorem 9.12. Suppose X is an almost platonic complex with surface structure. Then X is orientable. If X is finite, then X is not simply-connected.

Proof. Since X has surface structure, every edge has two incident faces. By (2) and (3) of Proposition 9.2, the configuration of two fev orbits are as shown in Figure 9.10. We can define the orientation of a face at an edge as from black to white, and this well defines the orientation of each face. Such orientations of faces are compatible at every edge, and therefore X is orientable. Now suppose that X is finite. If X is simply-connected, then X satisfies all conditions of Proposition 2.1. By the proposition, X must be platonic, violating our assumption. Therefore X can not be simply-connected. $\square$

Remark. It seems that simple connectedness will result in extra symmetries, at least for complexes with surface structure. Hence we doubt the existence of finite simply-connected almost platonic complexes. As for infinite cases, we seem to have more flexibility, but so far we have not found an example. A natural thought is to consider $\mathrm{Cox}_m\Gamma$ for a half-transitive graph Γ. The resulting complex is not platonic due to the restriction of its link. However it is not almost platonic either, since a generator of order 2 can reflect a face at an edge. It is of the type illustrated on the right of Figure 9.1.

The theory of the symmetry of complex tensor products can help us to obtain more almost platonic complexes, and the resulting complexes are mostly not surfaces.

Theorem 9.13. Suppose that X is an almost platonic simple prime complex, and Y is a platonic simple complex. Moreover, suppose that the 1-skeletons of X and Y are finite simple connected non-bipartite R-thin graphs with more than one vertex. Then the complex tensor product $X \otimes Y$ is an almost platonic complex.

Proof. Note that X^1 and Y^1 are symmetric graphs. By the argument in Corollary 7.14, we know that $X \otimes Y$ is a simple complex with finite simple connected non-bipartite edge-transitive R-thin 1-skeleton with more than one vertex. By Theorems 7.11 and 7.12, $X \otimes Y$ has a prime factorization $X \otimes Y_1 \otimes Y_2 \otimes \ldots \otimes Y_n$, where $Y = Y_1 \otimes Y_2 \otimes \ldots \otimes Y_n$, and $\mathrm{Aut}(X \otimes Y)$ is generated by automorphisms of prime factors and permutations of isomorphic factors. By Corollary 7.13, each Y_i is platonic, and hence not isomorphic to the almost platonic X. This implies that $\mathrm{Aut}(X \otimes Y) = \mathrm{Aut}(X) \times \mathrm{Aut}(Y)$.

To show that $X \otimes Y$ is almost platonic, we verify the conditions in Proposition 9.2. The *ev*-transitivity of $X \otimes Y$ follows easily from Definition 5.7. For f-transitivity, let f and f' are two arbitrary faces of $X \otimes Y$. Suppose that f is generated by combining a corner c_1 of a face f_1 in X with a corner c_2 of a face f_2 in Y, whereas f' is generated by combining a corner c_1' of a face f_1' in X with a corner c_2' of a face f_2' in Y. Since X is almost platonic, there exists $\rho \in \mathrm{Aut}(X)$ mapping c_1 to c_1'. Note that f' are generated by combining c_1' with c_2' in a particular orientation. Since Y is platonic, $\mathrm{Aut}(Y)$ can map c_2 to c_2' in either orientation, and there exists $\sigma \in \mathrm{Aut}(X)$ mapping c_2 to c_2' such that $(\rho, \sigma) \in \mathrm{Aut}(X) \times \mathrm{Aut}(Y)$ maps f to f'. Hence $X \otimes Y$ is f-transitive. Also there exists $\delta_1 \in \mathrm{Aut}(X)$ rotating f_1 by 1, and there exists $\delta_2 \in \mathrm{Aut}(Y)$ rotating f_2 by 1. Depending on how the corners c_1 and c_2 are combined to generate f, either (δ_1, δ_2) or $(\delta_1, \delta_2^{-1})$ can rotate f by 1.

We have shown (1) and (2) of Proposition 9.2. Suppose that (3) does not hold in $X \otimes Y$. In other words, there exists $\varphi \in \mathrm{Aut}(X \otimes Y)$ such that φ flips an edge e and stabilizes a face f incident to e. Suppose that e is generated by e_1 of X and e_2 of Y, and f is generated by f_1 of X and f_2 of Y. Note that $\mathrm{Aut}(X \otimes Y) = \mathrm{Aut}(X) \times \mathrm{Aut}(Y)$, and φ must be of the form $(\varphi_1, \varphi_2) \in \mathrm{Aut}(X) \times \mathrm{Aut}(Y)$. Then $\varphi_1 \in \mathrm{Aut}(X)$ flips e_1 and stabilizes f_1 incident to e_1, violating the almost platonicity of X. Hence $X \otimes Y$ satisfies (3), and is an almost platonic complex by Proposition 9.2. $\square$

Remark. In the proof of f-transitivity above, we do need the platonicity of Y to flip a face corner of Y. If X and Y are both almost platonic and $\mathrm{Aut}(X \otimes Y) = \mathrm{Aut}(X) \times \mathrm{Aut}(Y)$, then $X \otimes Y$ has two different face orbits, and is not almost platonic.

Theorem 9.13 gives a possible method to construct almost platonic complexes which are not surfaces. And now the question is, can we find an almost platonic complex X satisfying all the conditions in Theorem 9.13? The almost platonic simple complex T_1 in Figure 9.8 has 1-skeleton K_5, which meets all the conditions on the 1-skeleton. Is T_1 is a prime complex? Note that the link of T_1 is a cycle of length 4, which is obviously a prime graph. By Theorem 5.14, having a prime link implies that T_1 is a prime complex. Therefore T_1 satisfies all the conditions in Theorem 9.13.

What about X_Γ in Theorem 9.9, where Γ is a 1-transitive graph of valency 3? Note that the link of X_Γ is a cycle of length 3, which is clearly a prime graph, and hence X_Γ is a simple prime complex. Conder and Dobcsányi list all symmetric trivalent graphs up to 768 vertices in [7], but unfortunately there are only two 1-transitive non-bipartite graphs in the list, the F448A graph of girth 7 and the F504B graph of girth 9. Note that being R-thin results in a cycle of length 4, so these two graphs are not R-thin. Hence with these two graphs, we can construct X_Γ satisfying all the conditions in Theorem 9.13.

Consider the complex tensor product of the above X_Γ with a single polygon. By Theorem 9.13 the resulting complex is almost platonic, and by Theorem 5.14 the link of the product is the graph tensor product of a cycle of length 3 with K_2, which is a cycle of length 6. This means the complex tensor product is still a surface. To obtain almost platonic complexes which are not surfaces, we simply have to pick platonic complexes other than a single polygon to perform the complex tensor product.

Here our choice of 1-transitive Γ is greatly restricted by the non-bipartite condition. We want to point out that if Conjecture 8.11 is true, then there would be a similar theorem to Theorem 9.13 for complexes with bipartite 1-skeletons, and X_Γ from a 1-transitive bipartite Γ could also be used to generate almost platonic complexes.

Chapter 10

Computational Aspects

In this chapter we deal with computational problems about polygonal complexes. The first problem is about determining if two polygonal complexes are isomorphic. The second problem is about determining if we can build a complex with prescribed links. For these computation problems to make sense, all graphs and complexes discussed in this chapter are assumed to be finite.

The graph isomorphism problem is one of the few examples which are in the class of NP, but not known to be in P or NP-complete [14]. In fact, the complexity of the graph isomorphism problem is so special that computer scientists describe its polynomially equivalent problems as GI-complete problems. Note that the class of polygonal complexes contains the class of simple graphs, and therefore the polygonal complex isomorphism problem is at least as hard as the graph isomorphism problem. Now the question is, is it indeed theoretically harder?

The initial idea is to study a polygonal complex through the 1-skeleton of its barycentric subdivision, which carries extra information about faces. Nevertheless, we encounter some trouble immediately. For example, the barycentric subdivisions of a cube and an octahedron are isomorphic, not to mention the 1-skeletons of them. In fact if two polygonal complexes are surfaces dual to each other, then their barycentric subdivisions will be isomorphic. While such a phenomenon is of theoretical interest on its own, we will circumvent this obstacle by taking finer subdivisions.

For a polygonal complex X, we define the n-**subdivision graph** $\Gamma_n(X)$ of X as follows. First, we subdivide each edge of X by adding n intermediate vertices, which we call edge vertices. Secondly, add a vertex, which we call face vertex, to the middle of each face of X, and then add edges to join every face vertex to every vertex on the boundary

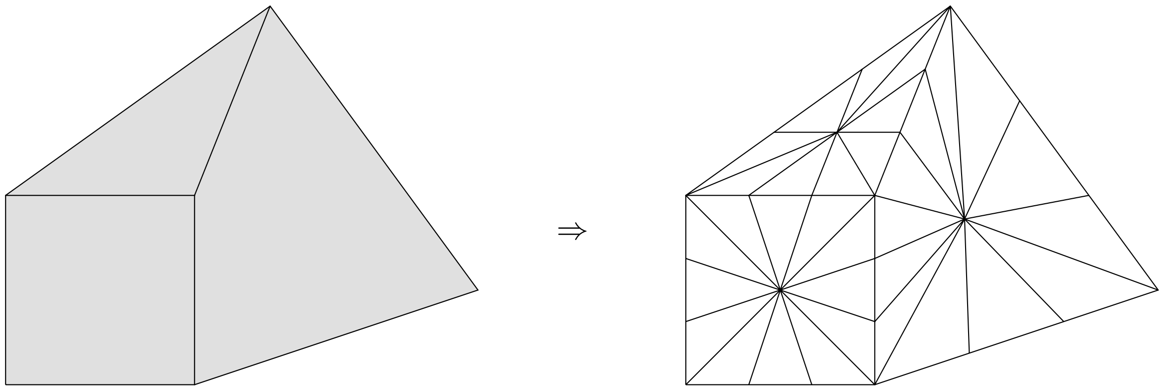

Figure 10.1: 2-subdivision graph of a complex

of the face. Finally, we take the 1-skeleton of the resulting complex to be $\Gamma_n(X)$. Vertices in $\Gamma_n(X)$ originating from vertices in $\Gamma(X)$ are called genuine vertices. When $n = 1$, $\Gamma_n(X)$ is the same as the 1-skeleton of the barycentric subdivision of X. Figure 10.1 is an example of 2-subdivision graph. Note that when we increase n for $\Gamma_n(X)$, only face vertices increase their valencies. Therefore if we choose large enough n, face vertices are those vertices of highest valencies.

For computational problems to make sense, we need to specify how to store polygonal complexes in a computer. A practical way to store a graph is to keep a list of vertices, and treat edges as pairs of vertices. To store a polygonal complex, in addition to the 1-skeleton, we treat faces as sequences of vertices according to the attaching maps.

Theorem 10.1. The polygonal complex isomorphism problem is polynomially equivalent to the graph isomorphism problem.

Proof. It suffices to show that we can polynomially reduce the polygonal complex isomorphism problem to the graph isomorphism problem. Suppose X and Y are two polygonal complexes. The goal is to show that X and Y are isomorphic if and only if $\Gamma_n(X)$ and $\Gamma_n(Y)$ are isomorphic for some n. The only if part is obvious from the definition. Note that in $\Gamma_n(X)$, the valency of a genuine vertex is the number of edges and faces incident to its original vertex in X, and the valency of an edge vertex is 2 plus the number of faces incident to the corresponding edge in X. We use $m(X)$ to denote the maximum of these numbers, which depends only on X, independent of n.

Take n to be the smallest integer such that $3(n+1) > \max\{m(X), m(Y)\}$, and assume that $\Gamma_n(X)$ and $\Gamma_n(Y)$ are isomorphic. First we try to recover X from $\Gamma_n(X)$. Note that in $\Gamma_n(X)$, a vertex is a face vertex if and only if its valency is at least $3(n+1)$, therefore we can distinguish face vertices easily. Now we remove all face vertices and incident edges from $\Gamma_n(X)$ to obtain a new graph $\Gamma_n^1(X)$. If there is a vertex v of valency not 2 in $\Gamma_n^1(X)$, then obviously v is a genuine vertex. Note that every vertex of distance $n+1$ to v is also a genuine vertex, and we can recursively determine all genuine vertices. The information about face vertices, edge vertices, and genuine vertices allows us to recover X completely, and Y can be recovered in the same manner as well. It is easy to see that an isomorphism from $\Gamma_n(X)$ to $\Gamma_n(Y)$ determines an isomorphism from X to Y.

What if every vertex in $\Gamma_n^1(X)$ has valency 2? This implies that $\Gamma_n^1(X)$ is actually a cycle. In this case we can not distinguish genuine vertices from edge vertices. Nevertheless, we can conclude that X is either a cycle or a polygon depending on if there is a face vertex in $\Gamma_n(X)$. Since $\Gamma_n(X)$ and $\Gamma_n(Y)$ are isomorphic, we know that X are Y are either isomorphic cycles or isomorphic polygons.

The above discussion shows that X and Y are isomorphic if and only if $\Gamma_n(X)$ and $\Gamma_n(Y)$ are isomorphic for the smallest integer n such that $3(n+1) > \max\{m(X), m(Y)\}$. Note that in $\Gamma_n(X)$ and $\Gamma_n(Y)$ the number of new vertices and edges is bounded by a polynomial function of the size of X and Y respectively. Hence the polygonal complex isomorphism problem is polynomially equivalent to the graph isomorphism problem. $\square$

Now we move to the problem of complexes with prescribed links. The one dimensional version of this problem has a long history, and can be stated as follows. For a sequence of non-negative integers $(d_1, d_2, \ldots, d_n)$, can we construct a graph of which valencies are exactly these numbers? If we allow parallel edges in the construction, it is easy to see that the construction is possible if and only if the sum of these integers are even. When we allow only simple graphs, this problem becomes much more technical, and the result is known as Erdős-Gallai Theorem [12]: For non-increasing $(d_1, d_2, \ldots, d_n)$, we can construct a simple graph with these valencies if and only if the sum of the sequence is even, and

$$\sum_{i=1}^{k} d_i \leq k(k-1) + \sum_{i=k+1}^{n} \min(d_i, k) \text{ for all } 1 \leq k \leq n.$$

The complex version of this problem can be stated as follows. For a finite sequence of graphs, can we construct a complex of which links are exactly these graphs? First we try to find some necessary condition. Let L be the disjoint union of links of a complex X. Note that an edge in X corresponds to two vertices with the same valency in L. Therefore, for any integer k, the number of vertices in L with valency k must be even. It turns out this condition is sufficient for the construction.

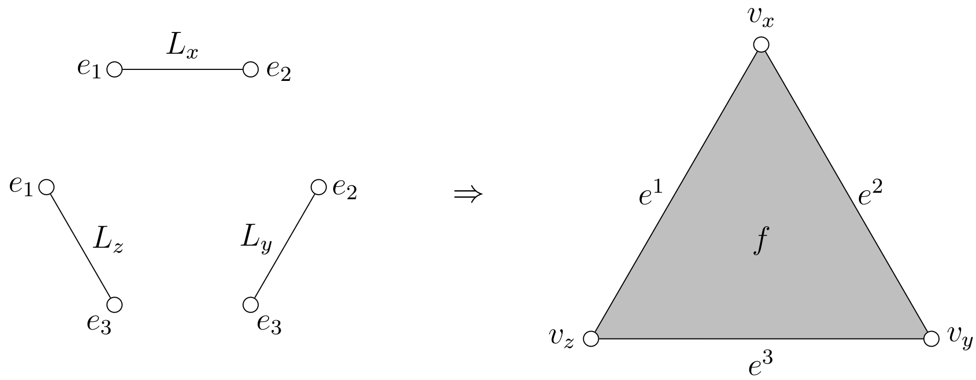

Figure 10.2: jump and walk

Proposition 10.2. Suppose that $(L_1, L_2, \ldots, L_n)$ is a sequence of graphs, and L is the disjoint union of these graphs. Then there exists a complex with $L_1, L_2, \ldots, L_n$ as its links if and only if for any integer k, the number of vertices in L with valency k is even.

Proof. The only if part is straightforward. For the if part, suppose that for any integer k, the number of vertices in L with valency k is even. Therefore in L, vertices with the same valency can be paired together without having any vertex left. We then assign a label to each pair, and label two vertices of the pair with such label. Suppose that $\{e_1, e_2, \ldots, e_m\}$ is the set of used labels. Note that each e_i appears exactly twice in L.

Now we build a complex X by tracing these labels. Let $\{v_1, v_2, \ldots, v_n\}$ be the vertex set of X, corresponding to $\{L_1, L_2, \ldots, L_n\}$. For each $1 \leq i \leq m$, find L_x and L_y with label e_i, and connect v_x and v_y by an edge e^i in X, as illustrated in Figure 10.2. Note that if e_i appears on some L_z twice, then we draw a loop e^i around v_z. This completes the 1-skeleton of X, and now we start to attach 2-cells. First pick an arbitrary edge in L, and suppose that this edge lies in L_x. We have the following two cases.

Case 1. The two vertices incident to this edge have the same label, say e_i. We attach a 2-cell along the loop e^i around v_x, and then mark this edge as used. Note that after this step, vertices in L with the same label still have the same unused valency.

Case 2. The two incident vertices have different labels, say e_1 and e_2. We mark this edge as used, and start from the e_2 vertex by taking two operations alternatingly. Operation one is to "jump" to another vertex with the same label, and record the label. Operation two is to "walk" to a neighbour through an unused edge, and then mark this edge as used. We stop the process when we fail the operation two. Note that when we jump between non-e_1 vertices, which have the same valency, the new vertex always has 1 more unused valency than the old vertex. When we jump from the initial e_1 vertex, in this case the new e_1 vertex has 2 more unused valency. Only when we jump to the initial e_1 vertex, we might run out of unused valency. Therefore when the process stops, we always record a sequence $(e_2, \dots, e_i, \dots, e_1)$ ending at e_1, which represents jumping back to L_x. We then attach a face starting at v_x along edges $(e^2, \dots, e^i, \dots, e^1)$, and this will bring us back to v_x, as illustrated in Figure 10.2. Note that vertices in L with the same label still have the same used valency.

In either case, vertices in L with the same label still have the same unused valency. If there is still an unused edge in L, we can apply the above method again to attach another face. After attaching finitely many faces, there will be no unused edge in L, and we finish the construction of X. Note the bijection between the corners of faces and the used edges in L. This shows that X has $L_1, \dots, L_n$ as links. $\square$

For this complex problem, we have a similar phenomenon as its graph counterpart. By imposing some obvious necessary condition, we can actually build a complex with prescribed links, although the proof is not as naive. Also like the graph version, imposing only obvious condition will allow examples with non-injective attaching maps. To construct a polygonal complex with prescribed links seems way too complicated to control. While the above construction does not guarantee a polygonal complex, it helps a lot to determine the 1-skeleton of the complex. We can determine the 1-skeleton simply by pairing vertices with the same valency, and the complex follows naturally. If we only ask to build a complex with a simple 1-skeleton, then we can actually forget about the graph structure of links, and only focus on valencies of each link.

Let $(L_1, L_2, \dots, L_n)$ be a sequence of graphs, d_i^j be the number of vertices of valency j in L_i, and m be the maximum valency among all L_i's. According to the proof of Proposition 10.2, to build a complex X of n vertices with link L_i at vertex v_i is equivalent to build a sequence of graphs $(\Gamma_1, \Gamma_2, \dots, \Gamma_m)$ with vertex set $\{v_1, v_2, \dots, v_n\}$ such that the valency of v_i in Γ_j is d_i^j. If we require X to have a simple 1-skeleton, it is equivalent to require each Γ_j is a simple graph, and $\Gamma_1, \Gamma_2, \dots \Gamma_m$ have at most one edge between any two vertices. This is related to the **simultaneous realization problem** of graphs: given two non-negative integral sequences $(d_1, d_2, \dots, d_n)$ and $(d_1', d_2', \dots, d_n')$, does there exist two simple graphs Γ and Γ' with vertex set $\{v_1, v_2, \dots, v_n\}$ such that the valencies of v_i in Γ and Γ' are d_i and d_i' respectively, and Γ and Γ' have at most one edge between any two vertices? Kundu [21] gives a necessary condition for the existence of such Γ and Γ', and

Chen [6] gives a short proof of Kundu's theorem which leads to a linear time constructive algorithm under Kundu's condition. But in general the simultaneous realization problem is difficult, as Dürr, Guiñez, and Matamala give the following result in [11].

Theorem 10.3. The simultaneous realization problem is NP-hard.

Corollary 10.4. Given a sequence $(L_1, L_2, \dots, L_n)$ of graphs, to determine if there exists a complex with $L_1, L_2, \dots, L_n$ links and having a simple 1-skeleton is NP-hard.

Proof. Suppose that we are given two non-negative integral sequences $(d_1, d_2, \dots, d_n)$ and $(d'_1, d'_2, \dots, d'_n)$. For each $i \in \{1, 2, \dots, n\}$, we try to construct a graph Γ_i with $d_i + d'_i$ vertices such that d_i vertices are of valency 2, and d'_i vertices are of valency 4. Since the valency sum is even, such construction is always possible. Now consider the sequence $(\Gamma_1, \Gamma_2, \dots, \Gamma_n)$ of graphs. From the discussion above, we know that to construct a complex with $\Gamma_1, \Gamma_2, \dots, \Gamma_n$ links and having a simple 1-skeleton is equivalent to the simultaneous realization problem for the two sequences $(d_1, d_2, \dots, d_n)$ and $(d'_1, d'_2, \dots, d'_n)$. Note that these two sequences can be given arbitrarily, and hence to construct such a complex has the same complexity as the simultaneous realization problem, which is NP-hard by Theorem 10.3. Now if we are given an arbitrary sequence $(L_1, L_2, \dots, L_n)$ of graphs, the construction is at least as hard as the simultaneous realization problem, and therefore is NP-hard as well. $\square$

Remark. In fact we can construct connected Γ_i easily, and hence this problem remains NP-hard when restricted to sequences of connected graphs.

Appendix

GAP Programs

Program 1.

We define a function Link(n) which outputs graphs of n vertices satisfying conditions (1) and (2) in Chapter 2. The value of n is restricted between 2 and 30.

```
LoadPackage("Grape");
Link:= function(n)
local list, i, G, Gn, a, b, li, N, m, edge, x, gamma;
list:=[];
for i in [1..NrTransitiveGroups(n)] do
   G:=TransitiveGroup(n,i);
   Gn:=Stabilizer(G,n);
   a:=Size(G);
   b:=a/n;                                        #size of Gn
   li:=Orbits(Gn,[1..n-1]);
   for N in li do
      m:=Length(N);
      if m>2 and n-1>m and b=Factorial(m) then
         edge:=[];
         for x in N do
            Add(edge,[x,n]);
            Add(edge,[n,x]);
         od;
         gamma:=EdgeOrbitsGraph(G,edge);
```

```
        if Size(AutGroupGraph(gamma))=a then
          Add(list,[i,N,not(IsomorphismGroups(Gn,SymmetricGroup(m))=fail),
          StructureDescription(G),gamma]);
        fi;
      fi;
    od;
od;
return list;
end;
```

Program 2.

This program verifies the $(14, 16)$-graph in Table 3.1 is isomorphic to the graph in Figure 3.3. This program also verifies that the $(28, 80)$-graph is a double cover of the $(14, 16)$-graph. The function Link(n) defined in Program 1 is used.

```
LoadPackage("Grape");
Link(14);
gamma1:=last[1][5];                               #the (14,16)-graph

g2:=Group((1,3,5,7,9,11,13)(2,4,6,8,10,12,14));
edge2:=[[14,3],[14,5],[14,7],[14,11],[3,14],[5,14],[7,14],[11,14]];
gamma2:=EdgeOrbitsGraph(g2,edge2);                #Figure 3.3

IsIsomorphicGraph(gamma1,gamma2);                 #true
GraphIsomorphism(gamma1,gamma2);                  #(2,4,8)(3,11,7)(5,9)(6,10,14)

Ceiling:=function(n)
if n=Int(n) then return Int(n);
else return Int(n+1); fi; end;                    #define ⌈n⌉

Link(28);
gamma3:=last[1][5];                               #the (28,80)-graph
```

```
edge3:=DirectedEdges(gamma3);
edge4:=[];
edge5:=[];

for i in edge3 do
   a:=[Ceiling(i[1]/2),Ceiling(i[2]/2)];
   if not (a in edge4) then
      Add(edge4,a);
   else
      Add(edge5,a);
   fi;
od;

Size(Set(edge5))=Size(edge5);                   #true, no duplicate in edge5
Set(edge4)=Set(edge5);                          #true, a double cover
g4:=Group((1,2)(3,4)(5,6)(7,8)(9,10)(11,12)(13,14));
gamma4:=EdgeOrbitsGraph(g4,edge4);              #quotient of gamma3

IsIsomorphicGraph(gamma1,gamma4);               #true

Link(20);
dodecahedron:=last[1][5];                       #dodecahedron graph
edgedodeca:=DirectedEdges(dodecahedron);
edgedouble:=[];

for i in edgedodeca do                          #part of edges in double cover
   Add(edgedouble,[i[1],i[2]+20]);
od;

g5:=Group((1,21)(2,22)(3,23)(4,24)(5,25)(6,26)(7,27)(8,28)(9,29)(10,30)(11,31)
          (12,32)(13,33)(14,34)(15,35)(16,36)(17,37)(18,38)(19,39)(20,40));
double:=EdgeOrbitsGraph(g5,edgedouble);         #double cover
AutGroupGraph(double);
Size(last);                                     #480
```

Program 3.

This program investigates how two faces intersect each other in Figure 9.2. Note that in GAP the multiplication of permutations is from left to right.

```
a_1:=(5,6)(3,4)(1,2);
b:=(6,8,9)(4,5,7)(1,2,3);
a_2:=(a_1)^b;
a_3:=(a_2)^b;

list1:=[()];
for i in [1..23] do
   if i mod 3 = 1 then Add(list1,(a_1)*list1[i]); fi;
   if i mod 3 = 2 then Add(list1,(a_2)*list1[i]); fi;
   if i mod 3 = 0 then Add(list1,(a_3)*list1[i]); fi;
od;

list2:=[(a_1)];
for i in [1..23] do
   if i mod 3 = 1 then Add(list2,(a_1)*list2[i]); fi;
   if i mod 3 = 2 then Add(list2,(a_2)*list2[i]); fi;
   if i mod 3 = 0 then Add(list2,(a_3)*list2[i]); fi;
od;

#list1=[ (), (1,2)(3,4)(5,6), (1,6,5,2,4,3)(7,8), (2,8,7,4,9,3,6,5),
         (1,8,7,4,6,2)(3,9), (1,5,8,4,6,2,9,3), (2,7,9,6)(3,5,8,4),
         (1,7,9,6,8,4,5,2), (1,2,3)(4,5,7)(6,8,9), (2,4,6,8,9,5,7,3),
         (1,4,2)(3,6,7)(5,8,9), (1,8,3)(2,6,7,9,5,4), (2,9)(3,8)(4,5)(6,7),
         (1,9,2)(3,5,7,6,4,8), (1,7,3)(2,5,9)(4,8,6), (2,3,7,5,9,8,6,4),
         (1,3,2)(4,7,5)(6,9,8), (1,4,7,6,9,8,5,3), (2,6,9,7)(3,4,8,5),
         (1,6,3,8,5,9,7,2), (1,9,7,5,6,3)(2,8), (2,5,6,3,9,4,7,8),
         (1,5,3,7,8,2)(4,9), (1,3)(2,7)(4,9) ]
#list2=[ (1,2)(3,4)(5,6), (), (1,5)(2,3)(7,8), (1,2,8,7,3,5)(4,9),
```

```
        (1,8,7,3,9,4,5,6), (1,6)(2,9,4,5,8,3), (1,2,7,9,5,8,3,6),
        (1,7,9,5)(3,4,6,8), (2,4,6,8,9,5,7,3), (1,2,3)(4,5,7)(6,8,9),
        (1,3,5,8,9,6,7,4), (1,8,4)(2,5,3)(6,7,9), (1,2,9)(3,8,4,6,7,5),
        (1,9)(3,6)(4,8)(5,7), (1,7,4,8,5,9)(2,6,3), (1,2,4)(3,7,6)(5,9,8),
        (1,4,7,6,9,8,5,3), (1,3,2)(4,7,5)(6,9,8), (1,2,5,4,8,6,9,7),
        (1,5,9,7)(3,8,6,4), (1,9,7,6,4,3,2,8), (1,2,6,4,7,8)(3,9),
        (1,6,5,4,9,3,7,8), (1,4,9,3,2,7)(5,6) ]
#Intersection(list1,list2)=[ (), (1,2)(3,4)(5,6), (1,2,3)(4,5,7)(6,8,9),
        (2,4,6,8,9,5,7,3), (1,3,2)(4,7,5)(6,9,8), (1,4,7,6,9,8,5,3) ]
```

Program 4.

This program lists all 1-transitive graphs of valency 3 up to 30 vertices.

```
LoadPackage("Grape");
list:=[];
for n in [2..30] do
   for i in [1..NrTransitiveGroups(n)] do
      a:=TransitiveGroup(n,i);
      b:=Stabilizer(a,n);
      li:=Orbits(b,[1..n-1]);
      for j in li do
         if Length(j)=3 and Size(Stabilizer(b,j[1]))=1 then
            edge:=[[n,j[1]],[j[1],n]];
            c:=EdgeOrbitsGraph(a,edge);
            if Size(AutGroupGraph(c))=Size(a) then
               Add(list,[n,i,j]);
            fi;
         fi;
      od;
   od;
od;
#list=[ [ 26, 6, [ 1, 17, 6 ] ], [ 26, 6, [ 4, 9, 11 ] ],
       [ 26, 6, [ 8, 20, 23 ] ], [ 26, 6, [ 13, 21, 15 ] ] ]
```

Program 5.

This program verifies 4 cases in Program 4 give isomorphic graphs, and Figure 9.3 and Figure 9.4 are two isomorphic drawing of them.

```
LoadPackage("Grape");
a:=TransitiveGroup(26,6);
b0:=EdgeOrbitsGraph(a,[[1,26],[26,1]]);            #w.r.t [1,17,6]
b1:=EdgeOrbitsGraph(a,[[4,26],[26,4]]);            #w.r.t [4,9,11]
b2:=EdgeOrbitsGraph(a,[[8,26],[26,8]]);            #w.r.t [8,20,23]
b3:=EdgeOrbitsGraph(a,[[13,26],[26,13]]);          #w.r.t [13,21,15]
IsIsomorphicGraph(b0,b1);                          #true
IsIsomorphicGraph(b0,b2);                          #true
IsIsomorphicGraph(b0,b3);                          #true

G1:=Group((1,2,3,4,5,6,7,8,9,10,11,12,13,14,15,
           16,17,18,19,20,21,22,23,24,25,26)^2);
edge1:=[[1,2],[2,1],[2,3],[3,2],[1,8],[2,21]];
fig1:=EdgeOrbitsGraph(G1,edge1);
IsIsomorphicGraph(b0,fig1);                        #true

G2:=Group((1,8,7,6,25,26)(2,9,14,5,24,19)(3,16,13,4,17,20)
          (21,10,15,12,23,18)(22,11));
edge2:=[[1,26],[26,1],[1,2],[2,1],[2,3],[3,2],[2,21],[21,2],
        [21,20],[20,21],[21,22],[22,21],[3,4],[4,3]];
fig2:=EdgeOrbitsGraph(G2,edge2);
Set(DirectedEdges(fig1))=Set(DirectedEdges(fig2));   #true
```

Bibliography

[1] L. Babai, W. Kantor, E. Luks, "Computational Complexity and the Classification of Finite Simple Groups", *Proc. 24th FOCS*, 162-171, 1983.

[2] W. Ballmann and M. Brin, "Polygonal Complexes and Combinatorial Group Theory", *Geom. Dedicata 50*, 165-191, 1994.

[3] N. Biggs, E. Lloyd, and R. Wilson, *Graph Theory 1736-1936*, Oxford University Press, 1999.

[4] B. Bollobas, *Extremal Graph Theory*, Dover Publications, 2004.

[5] M. Bridson and A. Haefliger, *Metric Spaces of Non-Positive Curvature*, Springer, 1999.

[6] Y. Chen, "A Short Proof of Kundu's k-Factor Theorem", *Discrete Math. 71*, 177-179, 1988.

[7] M. Conder and P. Dobcsányi, "Trivalent Symmetric Graphs Up to 768 Vertices", *J. Combin. Math. Combin. Comput. 40*, 41-63, 2002.

[8] H. S. M. Coxeter, *Regular Polytopes*, Dover Publications, Section 1.9, 1973.

[9] M. Davis, *The Geometry and Topology of Coxeter Groups*, Princeton University Press, 2008.

[10] W. Dörfler, "Primfaktorzerlegung und Automorphismen des Kardinalproduktes von Graphen", *Glasnik Mat. 9*, 15-27, 1974.

[11] C. Dürr, F. Guiñez, and M. Matamala, "Reconstructing 3-Colored Grids from Horizontal and Vertical Projections is NP-Hard: A Solution to the 2-Atom Problem in Discrete Tomography", *SIAM J. Discrete Math. 26*, 330-352, 2012.

[12] P. Erdős and T. Gallai, "Graphs with Prescribed Degrees of Vertices (in Hungarian)", *Matematikai Lapok 11*, 264-274, 1960.

[13] R. Frucht, "A One-Regular Graph of Degree Three", *Canad. J. Math. 4*, 240-247, 1952.

[14] M. R. Garey and D. S. Johnson, *Computers and Intractability: A Guide to the Theory of NP-Completeness*, W. H. Freeman, 1979.

[15] R. Hammack, W. Imrich, and S. Klavžar, *Handbook of Product Graphs*, second edition, CRC Press, 2011.

[16] F. Harary, *Graph Theory*, Addison-Wesley, pp. 173, 1994.

[17] A. Hatcher, *Algebraic Topology*, Cambridge University Press, 2002.

[18] D. Holt, "A graph which is edge transitive but not arc transitive", *Journal of Graph Theory 5*, 201-204, 1981.

[19] W. Imrich, "Automorphismen und das kartesische Produkt von Graphen", *Österreich. Akad. Wiss. Math.-Natur. Kl. S.-B. II 177*, 203-214, 1969.

[20] T. Januszkiewicz, I. J. Leary, R. Valle, and R. Vogeler, "Simple Platonic Polygonal Complexes", in preparation.

[21] S. Kundu, "The k-Factor conjecture is true", *Discrete Math. 6*, 367-376, 1973.

[22] R. Lyndon and P. Schupp, *Combinatorial Group Theory*, Springer, 2001.

[23] S. Mac Lane, *Categories for the Working Mathematician*, second edition, Springer, 1998.

[24] R. McKenzie, "Cardinal multiplication of structures with a reflexive relation", *Fund. Math. 70*, 59-101, 1971.

[25] D. J. Miller, "The automorphism group of a product of graphs", *Proc. Amer. Math. Soc. 25*, 24-28, 1970.

[26] J. R. Stallings, "Non-Positively Curved Triangles of Groups", *Group Theory from a Geometrical Viewpoint*, Trieste, 491-503, 1990.

[27] J. Świątkowski, "Trivalent Polygonal Complexes of Nonpositive Curvature and Platonic Symmetry", *Geom. Dedicata 70*, 87-110, 1998.

[28] J. Tits, "Le probléme des mots dans les groupes de Coxeter", *Symposia Mathematica (INDAM, Rome, 1967/68) Vol. 1*, Academic Press, London, 175-185, 1969.

[29] R. Valle, "Polygonal Complexes with Octahedral Links", *PhD thesis*, The Ohio State University, 2011.