COMBINATORIAL MAPS AND THE FOUNDATIONS OF TOPOLOGICAL GRAPH THEORY

A thesis presented in partial fulfilment of the requirements for the degree of Doctor of Philosophy in Mathematics at Massey University.

CRAIG PAUL BONNINGTON

1991

ABSTRACT

This work develops the foundations of topological graph theory with a unified approach using combinatorial maps. (A combinatorial map is an n-regular graph endowed with proper edge colouring in n colours.) We establish some new results and some generalisations of important theorems in topological graph theory. The classification of surfaces, the imbedding distribution of a graph, the maximum genus of a graph, and MacLane's test for graph planarity are given new treatments in terms of cubic combinatorial maps. Among our new results, we give combinatorial versions of the classical theorem of topology which states that the first Betti number of a surface is the maximum number of closed curves along which one can cut without dividing the surface up into two or more components. To conclude this thesis, we provide an introduction to the algebraic properties of combinatorial maps. The homology spaces and Euler characteristic are defined, and we show how they are related.

TABLE OF CONTENTS

PREFACE

Topological graph theory is concerned with the study of graphs imbedded in surfaces. During the past two decades, graph imbeddings in surfaces have received considerable analysis by combinatorial methods. The concept that motivated this thesis is the use of a special kind of edge coloured graph, called by Lins [11] a gem, to provide such a method to model graph imbeddings. For the most part, we shall push Lins' model further by using more general edge coloured graphs, called cubic combinatorial maps, to establish some new results and some generalisations of important theorems in topological graph theory. (A cubic combinatorial map is defined as a cubic graph endowed with a proper edge colouring in three colours.) It is the use of combinatorial maps that is the unifying feature in this thesis and its development of the foundations of topological graph theory.

An advantage of this approach over previous attempts to combinatorialise topological graph theory is that the theorems can be easily visualized, encouraging geometric intuition. We demonstrate how this axiomatic non-topological definition of a graph imbedding means that no topological apparatus needs to be brought into play when proving theorems in topological graph theory.

Following a chapter of introductory material, Chapter II gives a simple graph theoretic proof of the classification of surfaces in terms of cubic combinatorial maps. This provides our first example of the naturalness of cubic combinatorial maps as a variation on the simplicial complex approach to topology. As in topology, we can now assign an orientability character and genus or cross cap number to a given cubic combinatorial map. This chapter also serves as an introduction to the special operation or "move" on combinatorial maps that permeates this thesis.

In [24], Stahl presents a purely combinatorial form of the Jordan curve theorem from which graph theoretical versions (for example [26]) follow as corollaries. This was later (in [13]) presented in terms of cubic combinatorial maps. Generalisations of the Jordan curve theorem abound in topology, and therefore we make progress along these lines in Chapter III by presenting a generalisation of Stahl's work, motivated by the work in [13]. We give a combinatorial version of the theorem of topology which states that the first Betti number of a surface is the maximum number of closed curves along which one can cut without dividing the surface up into two or more components.

No text on the foundations of topological graph theory would be complete without some study of the set of surfaces a given graph can be imbedded on, or more precisely the imbedding distribution of a graph. The principal objective of topological graph theory is to determine the surface of smallest genus such that a given graph imbeds in that surface. In general, this surface is difficult to find. By way of contrast, the surface of largest genus such that a given graph imbeds in that surface can be found. We define a special partition of the set of all cubic combinatorial maps, and we say that two cubic combinatorial maps that belong to the same cell are congruent. In particular, two congruent gems correspond to two possible imbeddings for a given graph. We analyse the distribution of the genus or cross cap numbers associated with the cubic combinatorial maps congruent to a given one in Chapter IV. In Chapter V, we calculate the maximum value in this distribution. This work generalises results of Khomenko [9, 10] and Xuong [32].

In [30], short proofs of three graph theoretic versions of the Jordan curve theorem are given. In the spirit of Chapter III, we generalise the version, expressed in terms of a double cover for a graph, in Chapter VI. (A double cover is a family of circuits such that each edge belongs to exactly two.). Furthermore, we show how this work is related to our work on cubic combinatorial maps in Chapter III, and hence we proceed in the direction of Little and Vince in [14].

In an attempt to make a partial separation between graph theory and topology, MacLane proved that a given graph would be imbeddable on the sphere if and only if it had a certain combinatorial property. However, his characterization was proved by topological arguments. The tools introduced in Chapter VI which relate cubic combinatorial maps to double covered graphs are further applied in Chapter VII. Here we classify which cubic combinatorial maps are congruent to planar ones, where planarity is defined in terms of orientability and Euler characteristic. The classification given is a combinatorial generalisation of MacLane's test for planarity.

A more general version of the cubic combinatorial map is found by dropping the restriction of cubic graphs so as to include n-regular graphs. Of course we increase the number of colours for the edge colouring to n. To conclude this thesis, we provide an introduction to the algebraic properties of such maps. The homology spaces and Euler characteristic are defined, and we show how they are related. Furthermore, a general form of the "move" that permeates this thesis is presented, and we show how this move affects the Euler characteristic.

NOTES ON FIGURES

This thesis is mainly concerned with edge coloured graphs. Unfortunately, colour was not achieveable on laser printers at our disposal. It is possible, using the postscript language, to dash curved lines and to vary the width of a line. Therefore we represent the various colours by dashing edges according to the following figure.

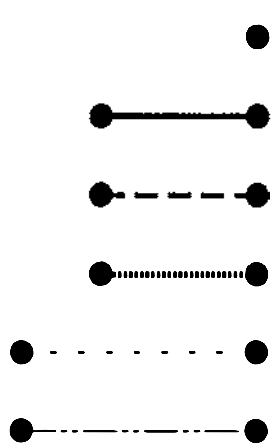

vertex

a blue edge

a red edge

a yellow edge

a path of red and yellow edges

a path of blue and yellow edges

For labellings, we will usually use a, b or c, together with a subscript or a prime, to label *red*, *blue* and *yellow* edges respectively. A vertex will always be labelled with u, v, w, x, y or z, together with a subscript or a prime.

ACKNOWLEDGEMENTS

A NUMBER OF PEOPLE HAVE ASSISTED ME IN MANY WAYS DURING MY DOCTORATE STUDIES.

To my supervisor, Charles Little, for his ready advice and unfailing encouragement, and for introducing me to the topic of my thesis.

To my colleagues and friends, Mark Byrne, Aroon Parshotam, John Giffin, Mike Hendy, Graeme Wake, Kee Teo, Ingrid Rinsma, Mike Steel and Mike Carter, for their comments, ideas, encouragement and friendship – particularly in the early stages of my doctorate studies.

To the University Grants Committee, for financial assistance.

DEDICATION

To

my wife, Karyn,

my Mum and Dad,

and my brother, Alex.

THIS THESIS WAS TYPESET ON APPLE LASERWRITER IINT PRINTERS IN TIMES 12/24, WITH HEADINGS IN COPPERPLATE AND PALATINO. DOCUMENT PREPARATION WAS ON APPLE MACINTOSH.COMPUTERS, USING MICROSOFT WORD, DESIGN SCIENCE MATH TYPE, ALDUS FREEHAND, AND CLARIS CAD. CROSS REFERENCING AND BIBLIOGRAPHY PREPARATION WERE ACHIEVED WITH CLARIS FILEMAKER PRO AND WORD REF.

Chapter I

INTRODUCTION

This chapter details the basic graph theoretical terminology which permeates this thesis. We shall state some standard results without proof.

1. Sets and Functions

Throughout this thesis, all sets are assumed finite. A set *contains* its elements and *includes* its subsets.

Let S and T be sets. Their *union*, $S \cup T$, and *intersection*, $S \cap T$, are defined in the usual way. These operations are commutative and associative. We denote by $S - T$ the collection of all elements of S not in T. The *sum*, $S + T$, of S and T is their symmetric difference, $(S - T) \cup (T - S)$, but we do not eschew the use of the latter notation. The operation of addition is also commutative and associative. If S is a non-empty collection of sets, their union, intersection and sum are denoted by $\bigcup S$, $\bigcap S$ and $\sum S$ respectively. We also define $\bigcup\varnothing = \sum\varnothing = \varnothing$. The cardinality of a set S is written as $|S|$. If S is a subset of T, we write $S \subseteq T$. We confine the use of the notation $S \subset T$ to the case where the inclusion is proper. $\mathbb{P}(T)$ denotes the collection of all subsets of a set T. For any non-negative integer n, $\mathbb{P}_n(T)$ denotes the collection of all $S \subseteq T$ for which $|S| \leq n$.

A *pair* is a set of cardinality 2. A special case is the pair $\{x, \{x, y\}\}$ for any objects x and y: this is the *ordered pair* or *ordered 2-tuple* (x, y), also denoted by xy or $(yx)^{-1}$. Its *components* are x and y.

Given objects $x_1, x_2, \ldots, x_n$ where $n > 2$, we define the *ordered n-tuple* $(x_1, x_2, \ldots, x_n)$, also denoted by $x_1 x_2 \ldots x_n$ or $(x_n x_{n-1} \ldots x_1)^{-1}$, to be $(x_1 x_2 \ldots x_{n-1})x_n$. Its *components* are $x_1, x_2, \ldots, x_n$. An *ordered 1-tuple* is a single object, its only *component*. The null set is the only *ordered 0-tuple* and has no components. An *ordered set* or *family* S is an ordered n-tuple for some non-negative integer n, and its cardinality $|S|$ is n.

We assume familiarity with the notion of a function and concomitant definitions. The domain of a function f is denoted Df. The restriction of f to a subset X of Df is denoted by $f|_X$. If f and g are functions with disjoint domains, then $f \cup g$ is the function h, with domain $Df \cup Dg$, such that $h|_{Df} = f$ and $h|_{Dg} = g$.

A *partition*, P, of a set X is a collection of non-empty disjoint subsets of X whose union is X. Thus if $X = \varnothing$ then $P = \varnothing$. The sets in P are the *cells* of the partition. If $Y \subseteq X$, then $P|_Y$ is the partition of Y whose cells are the non-empty intersections with Y of the cells of P.

2. GRAPHS

We define a *graph* G as an ordered triple $(VG, EG, \psi G)$ where VG and EG are sets and ψG is a function mapping EG into $\mathbb{P}_2(VG) - \{\varnothing\}$. We call VG, EG and ψG the *vertex set*, *edge set* and *incidence function*, respectively, of G. The elements of VG and EG are the *vertices* and *edges* respectively, of G. G is *null* if $VG = \varnothing$ and *empty* if $EG = \varnothing$. This thesis is concerned only with *finite* graphs, those graphs G in which the vertex set VG and edge set EG are both finite.

We say that a graph *contains* its vertices and edges. Accordingly we may describe the elements of $VG \cup EG$ as being *in* the graph G. The obvious extensions of this terminology to other graph theoretical concepts will be employed without being formally introduced.

Let G be a graph. An edge e is a *loop* or a *link* according to whether $|\psi G(e)| = 1$ or $|\psi G(e)| = 2$. The elements of $\psi G(e)$ are the *ends* of e, and e is *incident* on them and *joins* them. We write ψe for $\psi G(e)$ when no ambiguity emerges. (If v and w are the ends of e, we sometimes say that *e joins* $\{v, w\}$. This convention enables us to say that a link joins a pair of vertices.) We may also say that e joins v *to* w, or vice versa. The ends of e in turn are *incident* on e. Two distinct edges are *adjacent* (to each other) if they are incident on a common vertex. Two distinct vertices are *adjacent* (to each other) if they are incident on a common edge. Adjacent edges and adjacent vertices are sometimes described as *neighbours* (of each other).

The *degree*, $deg_G(v)$, in G of a vertex $v \in VG$ is $k + 2j$ where k and j, respectively, are the numbers of links and loops incident on v. We may delete the subscript in this notation if no ambiguity emerges. Vertices of degree 0 are *isolated*.

A graph is *n-regular* if the degree of every vertex is n. A 3-regular graph is *cubic*.

LEMMA I.1. *For any graph* G, $\sum_{v \in VG} deg(v) = 2|EG|$. ❑

A graph may be represented in the plane by a drawing in which distinct vertices are represented by distinct points. If v' and w' are points which represent distinct vertices v and w respectively, then an edge joining v and w is represented by a simple curve joining v' and w' and containing no other point which represents a vertex. If u' is a point representing a vertex u, then a loop incident on u is represented by a simple closed curve containing u' but no other point which represents a vertex. Thus an object which is both a vertex and an edge is represented twice, once by a point and once by a curve.

3. ISOMORPHISM OF GRAPHS

Graphs G and H are *isomorphic* if there exist bijections $f{:}VG \to VH$ and $g{:}EG \to EH$ such that $\psi H(g(e)) = f[\psi G(e)]$ for all $e \in EG$. We say that (f, g) is an *isomorphism* from G to H. Clearly an equivalence relation is herein defined. Usually one considers an equivalence class of isomorphic graphs rather than a single graph, since the internal structure of the vertices and edges escapes mention in the definition of a graph. It is therefore customary to identify the graphs in a particular equivalence class with a single representative selected from that class. Such a representative may be chosen to have its vertex set and edge set disjoint. Note that two isomorphic graphs may be represented by a common drawing.

4. SUBGRAPHS

A graph H is a *subgraph* of a graph G if $VH \subseteq VG$, $EH \subseteq EG$ and $\psi H = \psi G|_{EH}$. We say that a graph *includes* its subgraphs. For any $T \subseteq EG$, $G - T$ denotes the subgraph $(VG, EG - T, \psi G|_{EG - T})$, and $G[T]$ the subgraph $(W, T, \psi G|_T)$ where W is the set of all vertices incident on an edge of T.

Let H be a subgraph of G. H is a *proper* subgraph of G if $VH \cup EH \subset VG \cup EG$, and a *spanning* subgraph of G if $VH = VG$. A subgraph H of G which satisfies a specified property Z is a *minimal* subgraph of G satisfying Z if no proper subgraph of H satisfies Z. A subgraph H of G which satisfies a specified property Z is a *maximal* subgraph of G satisfying Z if H is not a proper subgraph of any other subgraph of G which satisfies Z.

5. COBOUNDARIES

We define the *coboundary*, $\partial_G v$, of a vertex v in a graph G as the set of all links incident on v. The *coboundary*, $\partial_G S$, of a set S of vertices is the sum of the coboundaries of those vertices. Thus $\partial_G\{v\} = \partial_G v$; $\partial_G v$ is a *vertex coboundary*. The symbol ∂_G may be replaced by ∂ if no ambiguity results. Clearly the edges of ∂S are those which join a vertex of S to a vertex of $VG - S$; hence $\partial S = \partial(VG - S)$. A set of edges is a *coboundary* if it is the coboundary of some set of vertices. A graph G is *connected* if $\partial S \neq \varnothing$ for every non-empty proper subset S of VG. A *component* of G is a maximal connected subgraph. Thus the components of a non-null graph are non-null. We denote by $c(G)$ the number of components in a graph G. Note that the word "component" has distinct meanings, but it should always be clear from the context which is intended.

An *isthmus* is the unique element of a coboundary of cardinality 1.

LEMMA I.2. *Let e be an isthmus in a connected graph G. Then $G - \{e\}$ has just two components.* ❑

Conversely, if G is a connected graph and e is an edge which is not an isthmus, then $G - \{e\}$ is connected, for otherwise $\{e\}$ would be a coboundary of G.

6. CIRCUITS, TREES AND PATHS

A *circuit* C in a graph G is the edge set of a minimal non-empty subgraph of G in which no edge is an isthmus. Its *length* is $|C|$. A circuit of length 1 has a loop as its unique element, and will sometimes be called a *loop*. Circuits of length 2, 3 and 4 are called *digons*, *triangles* and *squares* respectively.

LEMMA I.3. *Let G be a graph, and let C be a non-empty subset of EG. Then C is a circuit if and only if $G[C]$ is connected and each vertex of $VG[C]$ is of degree 2 in $G[C]$.* ❑

A graph is *bipartite* if every circuit has even length.

LEMMA I.4. *A graph is bipartite if and only if its edge set is a coboundary.* ❑

If G is a bipartite graph and S is a subset of VG such that $EG = \partial S$, then we say that $\{S, VG - S\}$ is a *bipartition* of G.

LEMMA I.5. *An edge is an isthmus if and only if it belongs to no circuit.* ❑

Let G be a graph. The *foundation*, fnd(G), of G is the subgraph spanned by the complement in EG of the set of isthmuses of G. Each edge of fnd(G) therefore belongs to a circuit, by Lemma I.5.

A *forest* is a graph in which every edge is an isthmus. A connected forest is a *tree*.

LEMMA I.6. *A connected graph is a tree if and only if it has no circuit.* ❑

LEMMA I.7. *A connected graph G is a tree if and only if $|EG| = |VG| - 1$.* ❑

A *path* P joining vertices v and w in a graph G is the edge set of a minimal connected subgraph containing v and w. We may also say that P *joins* v *to* w, or vice versa. Its *length* is $|P|$. We call v and w the *ends* of P. The other vertices of $VG[P]$ are said to be *internal vertices* of P.

LEMMA I.8. *A graph is connected if and only if every pair of vertices is joined by a path.* ❑

Let Q be a circuit or non-empty path in a graph G. We denote $VG[Q]$ by VQ. We say that Q *passes through* the elements of VQ, and we call VQ the *vertex set* of Q.

Let v, w be vertices in the same component of a graph G. By Lemma I.8, there exists a path joining v and w. The length of a shortest such path is the *distance* from v to w in G. If $v, w \in VD$, where D is a path or circuit in G, then the *distance* from v to w in D is the distance from v to w in $G[D]$.

Let v and w be vertices of a path P. Then we denote by $P[v, w]$ the edge set of the unique minimal connected subgraph of $G[P]$ containing v and w. $P[v, w]$ is a *subpath* of P.

Let v be a vertex of a circuit C, of length greater than 1 in a graph G. Let a, b be the edges of C incident on v. Then $C - \{a, b\}$ is also a path. We denote this path by C_v.

7. CONTRACTION

Let G be a graph and $\mathcal{Q}$ a partition of VG such that $G[Q]$ is connected for each $Q \in \mathcal{Q}$. Let H be a graph with vertex set $\mathcal{Q}$, and let EH be the set of edges of G which join vertices in distinct members of $\mathcal{Q}$. Suppose that each edge of H joins the two elements of $\mathcal{Q}$ containing its ends in G. Then H is called the *vertex contraction* of G determined by $\mathcal{Q}$. Let $S \subseteq EG$ and $T \subseteq EH$. We say that S *contracts* to T in H, or that T is a *contraction* of S in H, if H is a vertex contraction of G and $T = S \cap EH$. If $e \in EG, f \in EH$ and $\{e\}$ contracts to $\{f\}$ in H, then we say that e *contracts* to f or that f is a *contraction* of e.

8. CYCLE SPACES

A graph is *Eulerian* if every vertex has even degree.

LEMMA I.9. *A graph is Eulerian if and only if its edge set is a union of disjoint circuits.* ❑

Note first that if S is any set, then $\mathbb{P}(S)$ defines a vector space of dimension $|S|$ over the field of residue classes modulo 2, with the sum of sets defined in the usual way. We take $S = EG$ for some graph G, and consider some subspaces of $\mathbb{P}(EG)$.

LEMMA I.10. *For any graph G, the collection of the edge sets of the Eulerian subgraphs of G defines a vector space* $\mathbf{Z}(G)$. ❑

We call $\mathbf{Z}(G)$ the *cycle space* of G. Its elements are *cycles*. The orthogonal complement, $\mathbf{Z}^{\perp}(G)$, of $\mathbf{Z}(G)$ is the *cocycle space* of G.

LEMMA I.11. *For all graphs G, dim* $\mathbf{Z}(G) = c(G) - |EG| + |VG|$. ❑

9. 3-GRAPHS

Let K be a graph. A *proper edge colouring* of K is a colouring of the edges so that adjacent edges receive distinct colours. A 3-*graph* or *cubic combinatorial map* is defined as an ordered triple $(K, \boldsymbol{P}, \boldsymbol{O})$ where K is a cubic graph endowed with a proper edge colouring $\boldsymbol{P}$ in three colours and $\boldsymbol{O}$ is a ordering of the three colours. We shall assume throughout that the three colours are red, yellow and blue. We write $K = (K, \boldsymbol{P}, \boldsymbol{O})$ when no ambiguity results.

The set obtained from EK by deletion of the edges of a specified colour is the union of a set of disjoint circuits, called *bigons*. Thus bigons are of three types: red-yellow, red-blue and blue-yellow. We use denote the sets of red-yellow, red-blue and blue-yellow bigons by $\boldsymbol{B}(K)$, $\boldsymbol{Y}(K)$ and $\boldsymbol{R}(K)$ respectively. The total number of bigons is $r(K) = |\boldsymbol{B}(K)| + |\boldsymbol{Y}(K)| + |\boldsymbol{R}(K)|$.

EXAMPLE I.12. *Consider the* 3-*graph* K *in Figure* I.1. *Evidently* $\boldsymbol{B}(K) = \{\{a_1, c_2, a_3, c_4, a_4, c_3, a_2, c_1\}\}$, $\boldsymbol{Y}(K) = \{\{a_1, b_2, a_4, b_1\}, \{a_2, b_3, a_3, b_4\}\}$, *and* $\boldsymbol{R}(K) = \{\{c_1, b_1, c_3, b_4, c_4, b_2, c_2, b_3\}\}$. *Hence* $r(K) = 4$. ❑

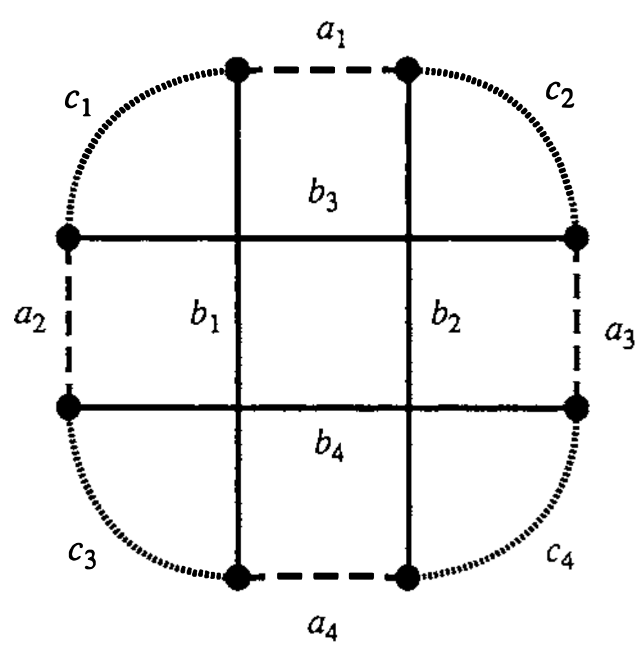

FIGURE I.1

10. GEMS

Following Lins[11, 12], we define a *gem* to be a 3-graph in which the red-blue bigons are quadrilaterals (circuits of length 4). For example, the 3-graph in Figure I.1 is a gem. We say that red-blue bigons in a gem are red-blue *bisquares*.

A 2-cell imbedding of a graph G, which may have loops, in a closed surface $\mathbb{S}$ can be modelled by means of a gem in the following way (see [11, 12, 14]). First construct the barycentric subdivision Δ of the imbedding of G, and colour each vertex of Δ with blue, yellow or red according to whether it represents a vertex, edge or face of the imbedding. Each edge of Δ then joins vertices of distinct colours, and may be coloured with the third colour. Let K be the dual graph of Δ, each edge of K being coloured with the colour of the corresponding edge of Δ. Then each red-blue bigon of the 3-graph K is a quadrilateral, so that K is a gem. (See Figure I.2. In this figure, the vertices of G are the solid circles and the edges are the thin solid lines joining such circles. All the circles are vertices of Δ. The edges of Δ are thin solid line segments and the thin dashed lines. The edges of K are thicker and coloured as indicated in the figure. The vertices of K should be self-evident.)

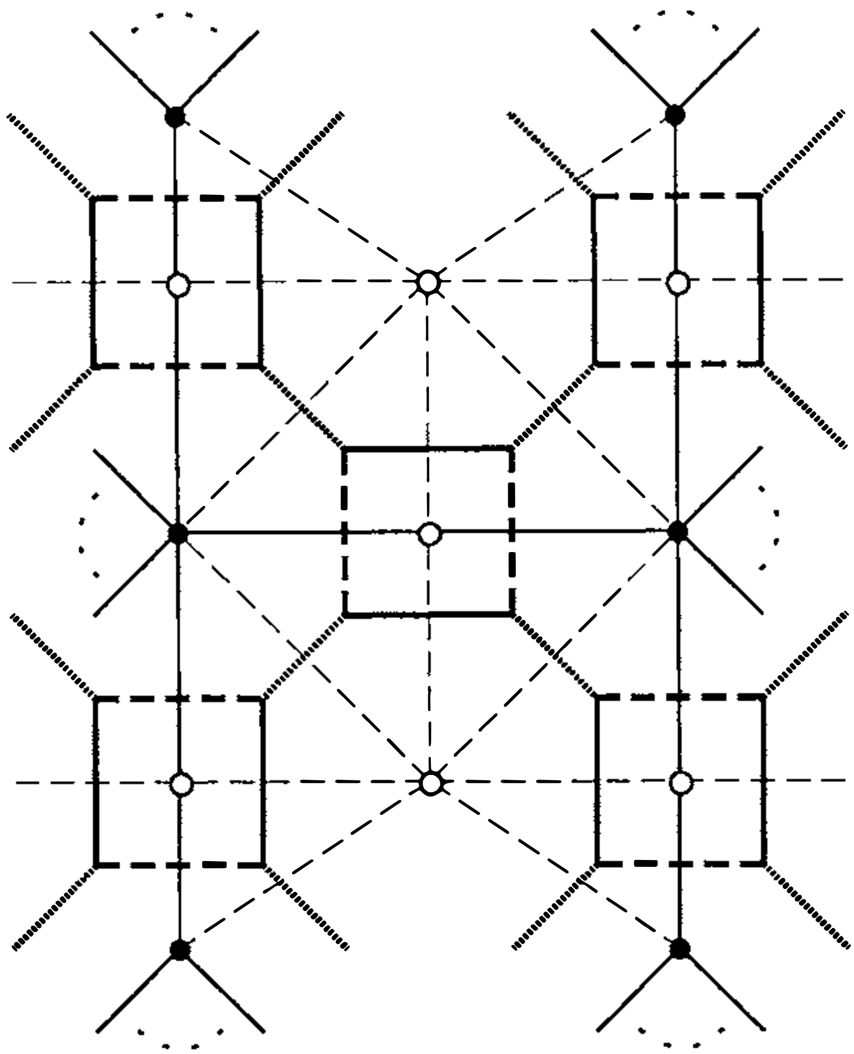

FIGURE I.2

This construction can be reversed. Given K, we first contract each red-yellow bigon to a single vertex. Each red-blue bigon then becomes a digon whose edges are both blue. The identification of the two edges in each of these digons yields G. Thus there is a 1:1 correspondence between gems and 2-cell imbeddings of graphs in closed surfaces. Also, if $S \subseteq EK$ then the blue edges in S appear in G. The set T of such edges of G is also said to *correspond* to S, and vice versa, but this correspondence is not 1:1. In general, T

corresponds to several subsets of EK. We also say that each of these subsets *represents* T.

If G is obtained from a gem K in this way, we say that G *underlies* K. We also say that K *represents* the imbedding of $G(K)$.

Lins also shows that the surface $\mathbb{S}$ is orientable if and only if K is bipartite. A generalisation of this result appears in [29].

The vertices of G are in 1:1 correspondence with the red-yellow bigons of K, the edges of G with the red-blue bigons of K, and the faces of the imbedding of G with the blue-yellow bigons of K. Thus we have $r(K) = |VG| + |EG| + |FG|$, where FG is the set of faces of the imbedding of G. The Euler characteristic $\chi(\mathbb{S})$ of $\mathbb{S}$ is therefore

$$\begin{aligned} |VG| - |EG| + |FG| &= r(K) - 2|EG| \\ &= r(K) - \frac{|VK|}{2} \end{aligned}$$

since $|VK| = 4|EG|$.

Gems appeared first in the doctoral dissertation of Robertson [20] and subsequently in work of Ferri and Gagliardi [4]. The correspondence between gems and imbeddings was developed by Lins in [11, 12], though his account was not expressed in terms of the barycentric subdivision of the imbedding.

Although it is gems that correspond to imbeddings, in this thesis we work in the more general setting of 3-graphs. The topological graph theory implications of this thesis are discovered by specialising the main theorems to the case of gems.

Chapter II

THE CLASSIFICATION OF COMBINATORIAL SURFACES

1. INTRODUCTION

In [27], Tutte approaches topological graph theory from a combinatorial viewpoint. In particular, an entirely combinatorial approach to the classification of surfaces is given. He uses the idea of a premap, which is expressed in [13] as a gem. Gems are also studied in [4, 28, 29] in the more general setting of n-graphs or combinatorial maps, a variation of the traditional simplicial complex approach to algebraic topology. In fact, the classification of surfaces in terms of 3-graphs is a direct consequence of the main theorem in [4], and is explicitly stated in [30]. Our purpose here is to show how the classification of surfaces by means of 3-graphs follows from

Tutte's approach in [27] and the relationship between 3-graphs and premaps. This approach provides a possible tool for proving theorems about cubic graphs with a proper edge colouring in three colours.

2. PREMAPS

The relationship between gems and premaps is established in [13]. Let X be a set such that $|X|$ is divisible by 4. Let θ and ϕ be permutations on X satisfying the conditions $\theta^2 = \phi^2 = I$ and $\theta\phi = \phi\theta$, and suppose x, θx, ϕx, $\theta\phi x$ are distinct for each $x \in X$. Let P be another permutation on X such that $P\theta = \theta P^{-1}$, and for each x let the orbits of P through x and θx be distinct. Then (X, θ, ϕ, P) defines a *premap*, M. Tutte also defines Ψ_L as the permutation group generated by a non-empty set L of permutations of X. Then the premap M is a *map* if for each $x \in X$ and $y \in X$ there is a permutation $\pi \in \Psi_{\{\theta, \phi, P\}}$ such that $\pi x = y$.

We now show how to construct a gem $K(M)$ that *represents* a premap M. Each element of X is represented by a vertex. For each $x \in X$, let us draw a red edge joining x to θx, a blue edge joining x to ϕx and a yellow edge joining x to $P\theta x$. (See Figure II.1.) It is shown in [13] that this construction yields a gem.

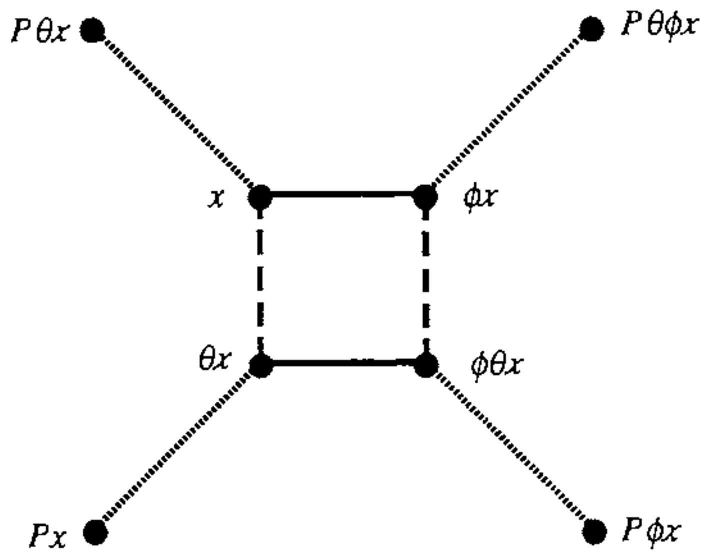

FIGURE II.1

Conversely for any gem K it is easy to construct a premap M for which $K = K(M)$. X is the vertex set of K, and the red and blue edges determine the involutions θ and ϕ respectively. Moreover, for any vertex x, Px is the unique vertex joined to θx by a yellow edge.

Evidently the premap M is a map if and only if $K(M)$ is connected.

Let M be a map. It is shown in [27, p.257] that the number of equivalence classes determined by $\Psi_{\{\theta\phi, P\}}$ is either 1 or 2, where two objects are regarded as equivalent if some permutation in $\Psi_{\{\theta\phi, P\}}$ maps one onto the other. We call these equivalence classes the *orientation classes* of M. The premap M is *orientable* if the number of orientation classes is 2, and *non-orientable* otherwise. The following lemma is proved in [12, 13] and a generalisation of it appears in [28, 29].

LEMMA II.1. *A map M is orientable if and only if $K(M)$ is bipartite.*

In general, we say a 3-graph is *orientable* if it is bipartite, and *non-orientable* otherwise.

We define the *Euler characteristic* of K to be

$$\begin{aligned} \chi(K) &= r(K) - |EK| + |VK| \\ &= r(K) - \frac{|VK|}{2} \end{aligned}$$

since K is cubic. (See [12, 13].) The *Euler characteristic* of a map M is $\chi(K(M))$.

Tutte [27] defines a surface as the class of all maps with a given Euler characteristic and given orientability character, provided that the class is non-empty. In our setting, such a class corresponds to a class of connected gems. However, we will work in the more general setting of 3-graphs and define a *surface* as the class of all connected 3-graphs with a given Euler characteristic and a given orientability character, provided that the class is non-empty. The main theorem of this chapter classifies all such surfaces. It states that one 3-graph can be obtained from another by a finite number of "moves" if and only if they belong to the same surface, where the moves are the crystallisation moves of [4] and are defined in the next section.

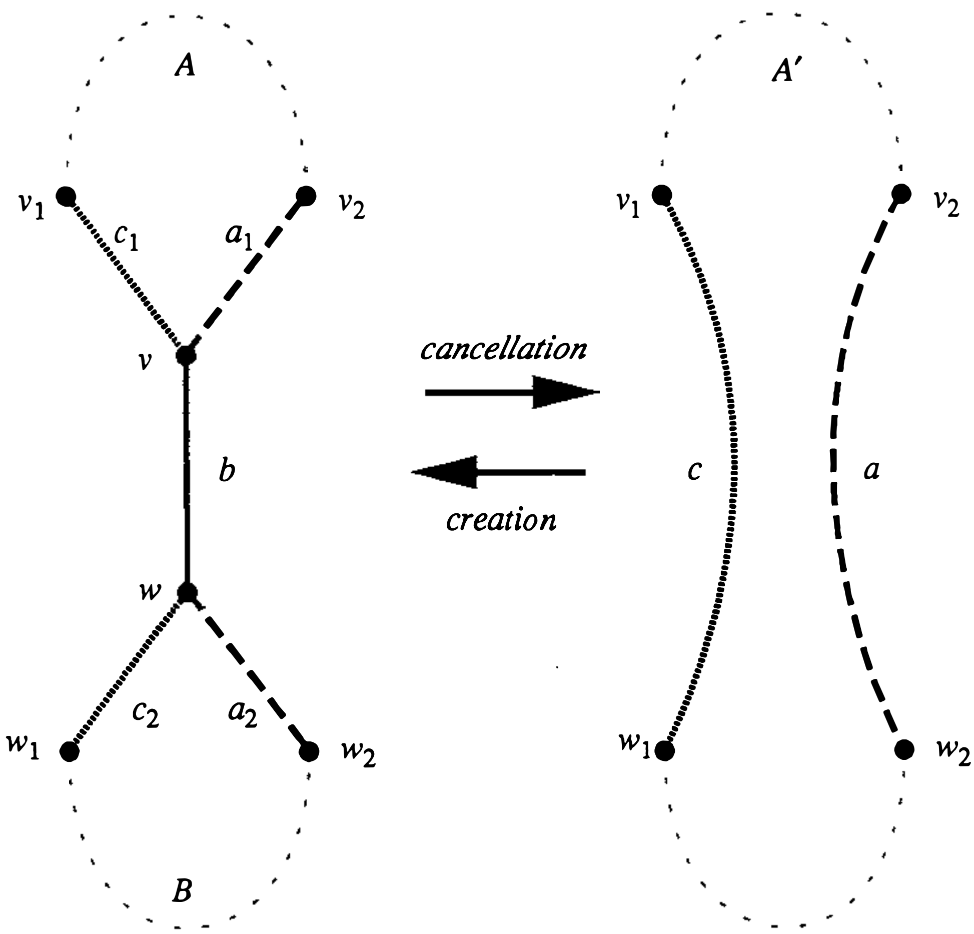

FIGURE II.2

3. DIPOLES

Let v and w be a pair of adjacent vertices in a 3-graph K. Suppose that v and w are linked by one edge b, which is blue. Following Ferri and Gagliardi [4], we say that b is a *blue* 1-*dipole* if the red-yellow bigons A and B passing through v and w respectively are distinct. Let c_1 and c_2 be the yellow edges incident on v and w respectively. Let a_1 and a_2 be the red edges incident on v and w respectively. Let v_1, v_2, w_1, w_2 be the vertices other than v and w incident on c_1, a_1, c_2, a_2 respectively. The *cancellation* of this blue 1-dipole b is the operation of deletion of the vertices v and w followed by the insertion of edges c and a linking v_1 to w_1 and v_2 to w_2 respectively. (See Figure II.2.) We denote the resulting 3-graph by $K - [b]$. We observe that A and B have coalesced into one red-yellow bigon A'. The *creation* of a blue 1-dipole is the inverse operation. Similar definitions can be made for red and yellow 1-dipoles.

Now suppose that v and w are linked by two edges a and b coloured red and blue respectively. Following Ferri and Gagliardi [4], we say that $\{a, b\}$ is a *red-blue 2-dipole* if the yellow edges c_1 and c_2 incident on v and w respectively are distinct. Let c_1 link v and v_1 and let c_2 link w and w_1. The *cancellation* of this red-blue 2-dipole is the operation of deletion of the vertices v and w followed by the insertion of an edge c linking v_1 to w_1. (See Figure II.3.) We denote the resulting 3-graph by $K - [a, b]$. We observe that c_1 and c_2 have coalesced into one yellow edge c. The *creation* of a red-blue 2-dipole is the inverse operation. Similar definitions can be made for red-yellow and blue-yellow 2-dipoles.

We note that the yellow edge c_1 is a yellow 1-dipole in K and that the 3-graph $K - [c_1]$ is isomorphic to $K - [a, b]$. Hence cancellation or creation of a 2-dipole is in fact a special case of a 1-dipole cancellation or creation.

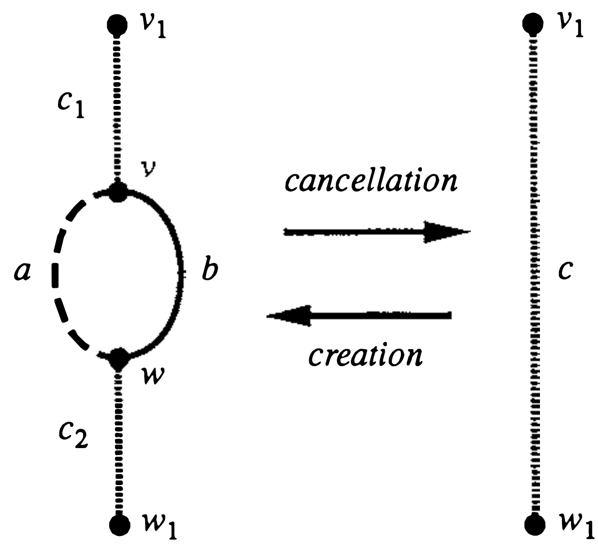

FIGURE II.3

A μ*-move* is a cancellation or creation of a 1-dipole. Two 3-graphs are μ*-equivalent* if one can be obtained from the other by a finite sequence of μ-moves. It is shown in [4] that two 3-graphs are equivalent if and only if the corresponding surfaces are homeomorphic. Thus the following theorem is equivalent to the classification of surfaces, due to Dehn and Heegaard [2]. Our proof is essentially theirs translated into the setting of coloured graphs. A similar proof in terms of premaps appears in [27]. The proof of the necessity appears as Theorem II.3, and the proof of the sufficiency appears in Sections 4 and 5.

THEOREM II.2. *Two connected* 3*-graphs* K *and* J *are* μ*-equivalent if and only if they have the same Euler characteristic and orientability character.*

THEOREM II.3. *If two connected 3-graphs K and J are μ-equivalent, then they belong to the same surface.*

Proof. We may assume that J is obtained from K by cancellation of a 1-dipole. We will show that $\chi(K) = \chi(J)$ and that K is bipartite if and only if J is bipartite. Indeed, in a 1-dipole cancellation the number of bigons drops by one and the number of vertices by two. Therefore

$$\begin{aligned}\chi(J) &= r(J) - \frac{|VJ|}{2} \\ &= r(K) - 1 - \frac{|VK|}{2} + 1 \\ &= \chi(K).\end{aligned}$$

Now assume that K is bipartite. Therefore, one may colour the vertices of K black or white so that adjacent vertices receive distinct colours. Evidently v and w receive distinct colours, as do v_1 and w_1, and v_2 and w_2. Hence we conclude that J is bipartite. Similarly K is bipartite if J is bipartite. ❑

We conclude this section by describing another operation on 3-graphs called cancellation of a red-blue bigon. This operation is in fact a pair of dipole cancellations.

Suppose Y to be a red-blue bigon of length 4 in a 3-graph K. Label the edges and vertices incident of Y as in Figure II.4a. If b is a blue 1-dipole then let $K' = K - [b]$. (See Figure II.4b.) Let a' denote the red edge of K' that links v' and w'. If $\{a', b'\}$ is a red-blue

2-dipole then let $K'' = K' - [a', b']$. (See Figure II.4c.) We say that K'' is obtained from K by *cancellation* of the red-blue bigon Y. Let c and c' denote the yellow edges that join v_1 and w_1, and v_2 and w_2, respectively. The inverse operation is described as *splitting* c and c' to *create* the red-blue bigon Y. By definition K and K'' are μ-equivalent.

a)

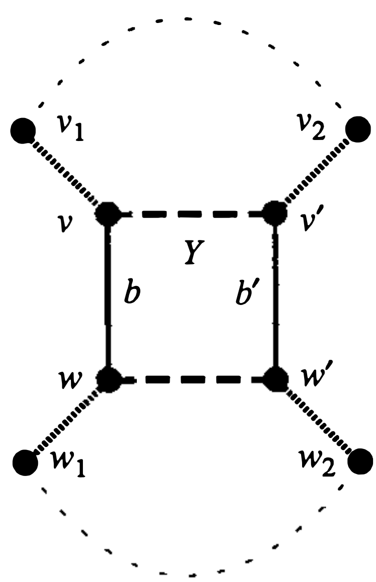

b)

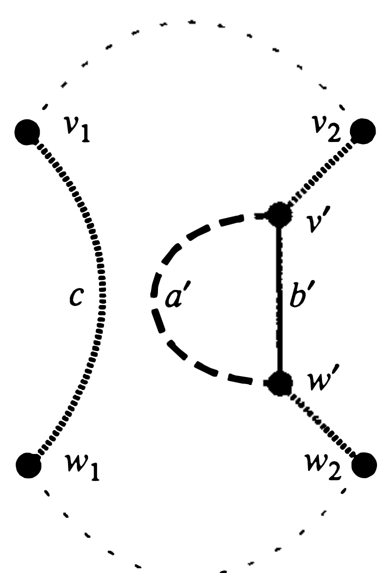

c)

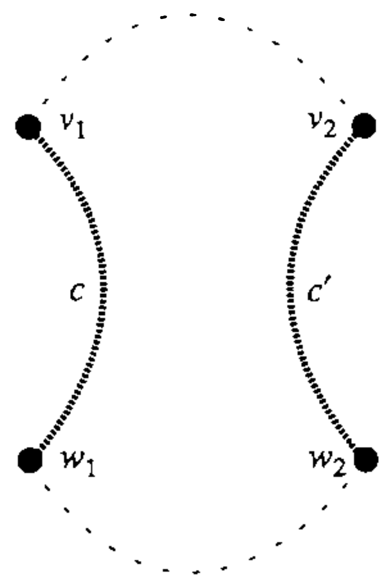

FIGURE II.4

4. REDUCED AND UNITARY 3-GRAPHS

The 3-graph with two vertices is *trivial.*

We assume given a connected 3-graph K. Suppose K has at least two red-yellow bigons. Since K is connected there must exist a blue 1-dipole in K. Cancelling this dipole reduces the number of red-yellow bigons by one. Proceeding inductively, we obtain a connected 3-graph which has exactly one red-yellow bigon. Similarly, we reduce to 1 the numbers of blue-yellow and red-blue bigons by red 1-dipole and yellow 1-dipole cancellations. The resulting 3-graph is a *reduced 3-graph* of K, and has just 3 bigons, one of each type. Note that a reduced 3-graph of K is μ-equivalent to K.

LEMMA II.4. *If K is a connected 3-graph then $\chi(K) \leq 2$. Moreover if $\chi(K) = 2$ then K is bipartite and the only reduced 3-graph of K is trivial.*

Proof. Let K' denote a reduced 3-graph of K. Since K and K' are μ-equivalent, we have $\chi(K) = \chi(K')$. Furthermore

$$\begin{aligned} \chi(K') &= r(K') - \frac{|VK'|}{2} \\ &\leq 2 \end{aligned}$$

since all 3-graphs have at least 2 vertices. If $\chi(K) = \chi(K') = 2$ then $|VK'| = 2$, and K' is the trivial 3-graph. Since the trivial 3-graph is bipartite, we have that K is bipartite by Lemma II.3. ❑

The *combinatorial sphere* is the class of all connected 3-graphs with Euler characteristic 2, and all 3-graphs in it are called *planar*. Thus the trivial 3-graph is planar.

A connected 3-graph K is *unitary* if $|B(K)| = |R(K)| = 1$. Hence reduced 3-graphs are unitary.

LEMMA II.5. *A connected 3-graph K is μ-equivalent to the trivial 3-graph or a unitary gem.*

Proof. Let K' denote a reduced 3-graph of K. If K is planar then K' is the trivial 3-graph by Lemma II.4 and we are done. Hence we assume otherwise. Let B denote the red-yellow bigon in K' and let Y be the red-blue bigon in K'. If Y is a digon (a circuit of length 2), then each yellow edge incident on Y is a yellow 1-dipole, a contradiction. If Y is a bisquare, then K' is a unitary gem and we are done. In the remaining case, let a_1 and a_2 be red edges of Y both adjacent to a common blue edge b. Let P_1 and P_2 be the two paths of $B - \{a_1, a_2\}$. Clearly there exist yellow edges $c_1 \in P_1$ and $c_2 \in P_2$, and therefore we split c_1 and c_2 to create a red-blue bisquare Y_1. Let K'' denote the resulting graph. Evidently b is a blue 1-dipole in K'', and therefore we cancel it to obtain a unitary 3-graph U. The red-blue bigon in U corresponding to Y has one blue edge less than Y. Proceeding inductively we obtain a unitary 3-graph U' where the red-blue bigon corresponding to Y is a bisquare. However U' is a unitary gem since all new red-blue bigons created were bisquares. ❑

Let us now study the case in which the surface is not the combinatorial sphere. Starting with an arbitrary 3-graph, we change it by μ-moves, as in Lemma II.5, into a unitary gem U.

Let B be the one red-yellow bigon in U. Let Y be an arbitrary red-blue bigon of U. Label the edges and vertices incident on Y as in Figure II.4a. If $w' \in V(B_{v'}[v, w])$ then we say that Y is a *cross-cap* of B. (See Figure II.5.) It is an *assembled* cross-cap of B if there exists a subpath of B with just two red edges, both in Y. If Y is not a cross-cap then it is a *cap* of B.

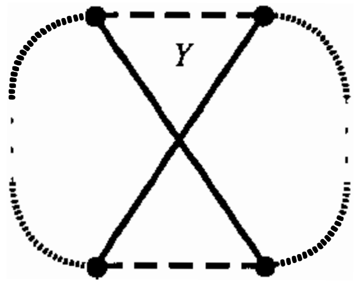

FIGURE II.5

Suppose there exist two caps X and Y of B. Again we label the edges and vertices incident on Y as in Figure II.4a. If $|X \cap B_{w'}[v, w]| = 1$ then we say that X and Y are *bound* in B, and $\{X, Y\}$ is a *handle* of B. (See Figure II.6.) Such a handle is *assembled* if there exists a subpath of B with just four red edges, all in $X \cup Y$.

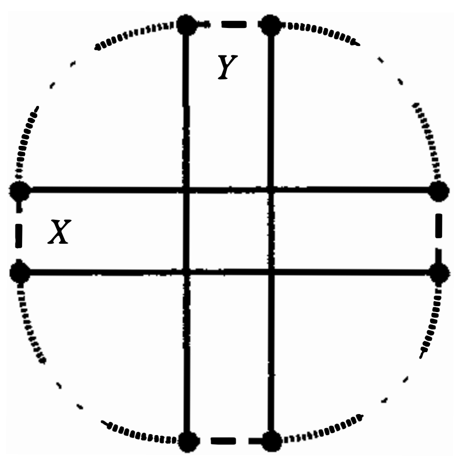

FIGURE II.6

LEMMA II.6. *A given unitary gem U is μ-equivalent to another unitary gem J for which all the cross-caps are assembled.*

Proof. Let B be the one red-yellow bigon in U and let Y be an unassembled cross-cap. Label the edges of B incident on vertices of

VY as in Figure II.7a. Split d and d' to create a red-blue bigon X and let U' denote the resulting gem. (Figure II.7b.) In U', a and a' belong to distinct red-yellow bigons. Hence we may let U'' denote the gem obtained by cancelling Y in U'. (Figure II.7c.) Clearly $|B(U'')| = |B(U)| = 1$ and $|R(U'')| = |R(U)| = 1$. It is also clear that X is an assembled cross-cap in U''. Moreover, no assembled cross-cap in U has been lost; the red edges of such a cross-cap are in $B - \{c, a, d, c', a', d'\}$.

By this procedure we can reduce to the case in which every cross-cap of U is assembled. ❑

a)

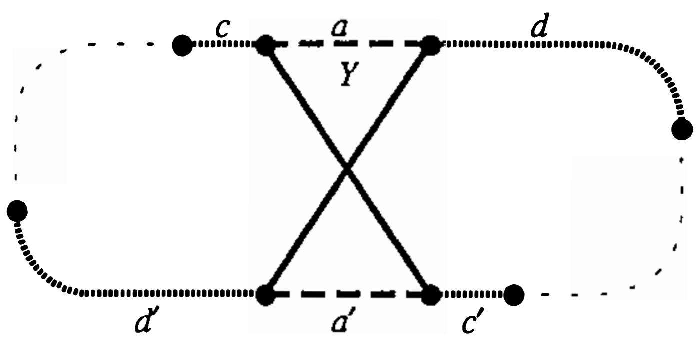

b)

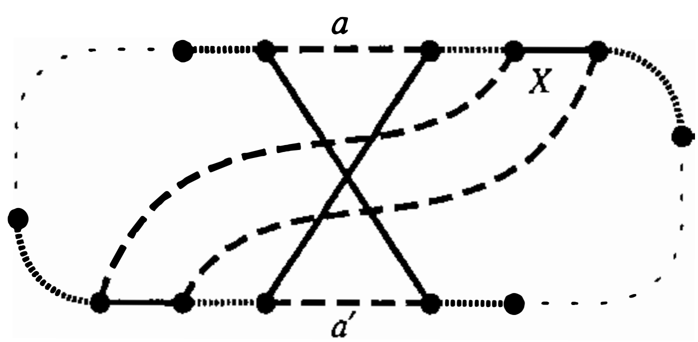

c)

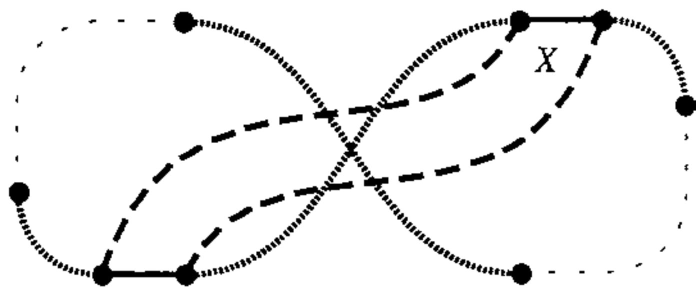

FIGURE II.7

LEMMA II.7. *Let U be a unitary gem with all cross-caps assembled. Then all caps in U are bound.*

Proof. Suppose there exists an unbound cap Y in U. Label the vertices and edges adjacent to Y as in Figure II.4a. Since there is no blue edge in U that has terminal vertices in both $VB_v[v', w']$ and $VB_{v'}[v, w]$, b and b' must belong to distinct blue-yellow bigons. This contradicts the fact that U is unitary. ❑

Consider our unitary gem in which all cross-caps are assembled. Then Lemma II.7 tells us that either the red-blue bigons are all assembled cross-caps or there are two red-blue bigons that constitute a handle. The next lemma deals with the assembly of handles.

LEMMA II.8. *By a finite sequence of μ-moves, we can convert a given unitary gem into one in which each red-blue bigon is an assembled cross-cap or a member of an assembled handle.*

Proof. Suppose there exists a handle $\{X, Y\}$ in U. The following uses the notation in Figure II.8a. Split d_2 and d_4 to create a red-blue bigon W and let U_1 denote the resulting gem. (See Figure II.8b.) In U_1, a_1 and a_3 are in distinct red-yellow bigons, and so the operation of cancellation of a red-blue bigon is applicable to X. We apply it, and let U_2 denote the resulting graph. (See Figure II.8c.) In U_2 let c_5 denote the yellow edge adjacent to a_4 other than c_4. Let b_5 denote the blue edge of W adjacent to c_5, and let d_5 denote the yellow edge other than c_5 adjacent to b_5. Split c_5 and d_5 to create a red-blue bigon Z and let U_3 denote the resulting gem. (See Figure II.8d.) In U_3, a_2 and a_4 are in distinct red-yellow bigons, and so the operation of cancellation of a red-blue bigon is applicable to Y. We apply it and let U' denote the resulting gem. (See Figure II.8e.)

The above process transforms U into another unitary gem U'. The handle $\{X, Y\}$ has been replaced by the assembled handle $\{W, Z\}$. Any assembled cross-cap or other assembled handle of B has edges in $B - \{a_1, c_1, d_1, a_2, c_2, d_2, a_3, c_3, d_3, a_4, c_4, d_4\}$, and is preserved.

By repetition of the operation just described, we can replace unassembled handles by assembled ones until we have a unitary gem of the kind required. (No red-blue bigons will be left over, by Lemma II.7.) ❑

a)

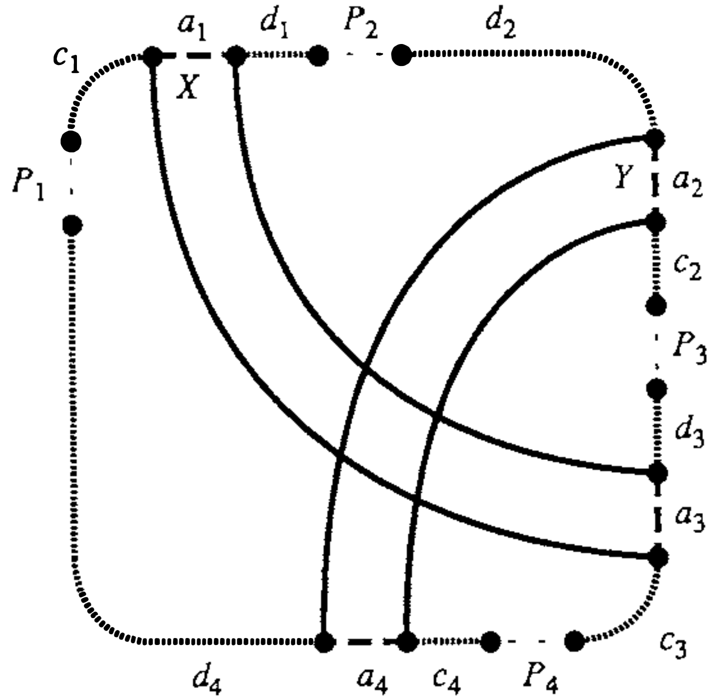

b)

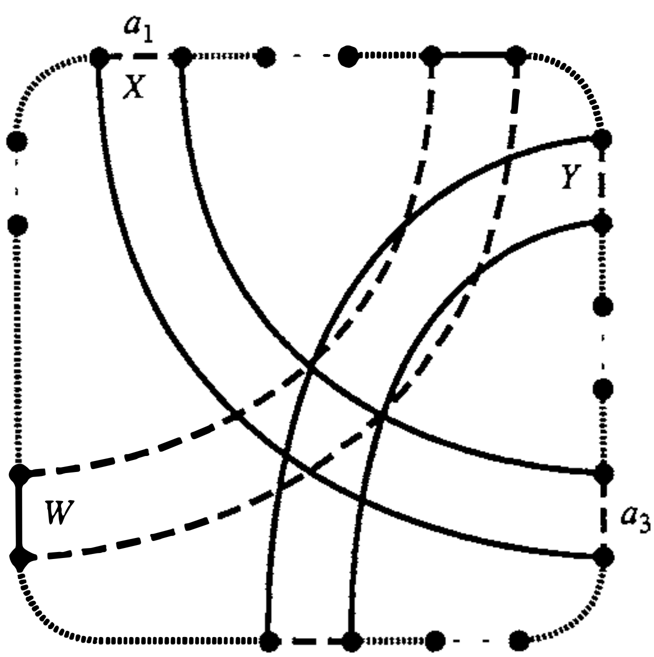

c)

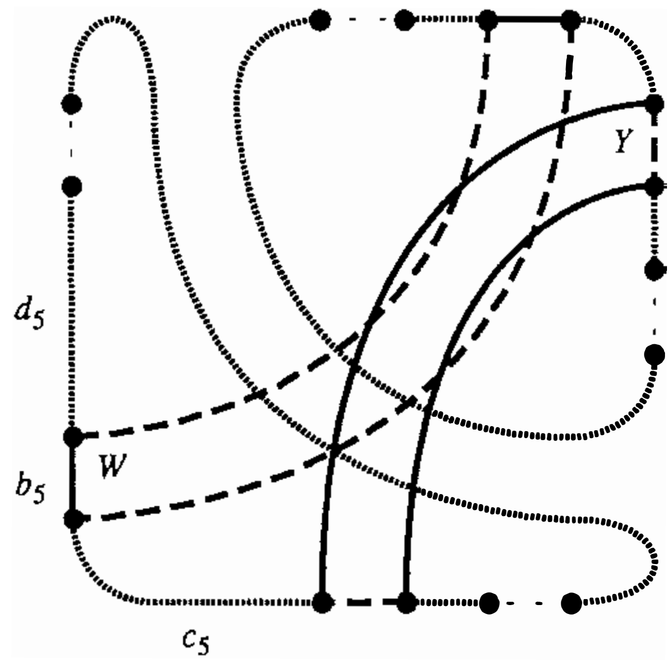

d)

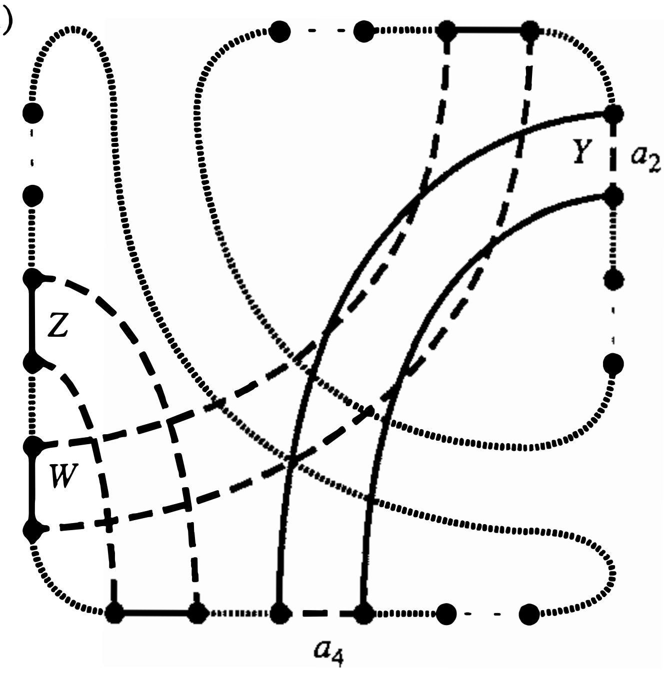

e)

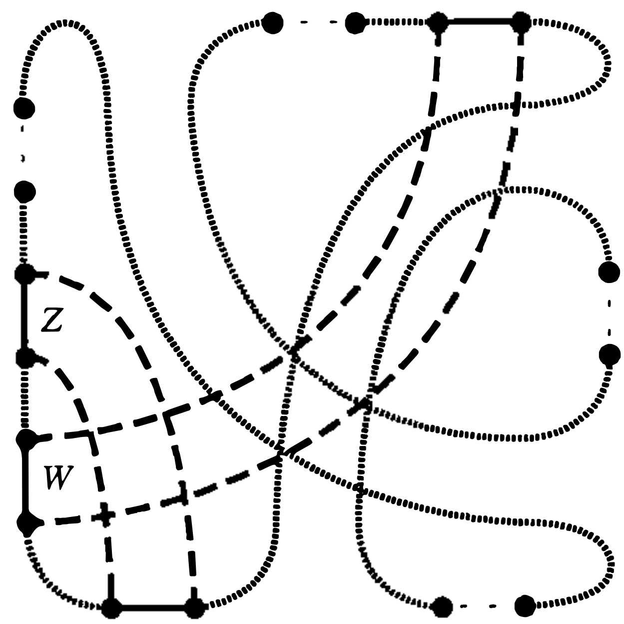

FIGURE II.8

LEMMA II.9. *By a finite sequence of μ-moves we can convert a given unitary gem into one in which all red-blue bigons are assembled cross-caps or all red-blue bigons are members of assembled handles.*

Proof. We may suppose our unitary gem U already in the form specified in Lemma II.8. Suppose it to have at least one assembled cross-cap and at least one assembled handle. Then the situation arising is depicted in Figure II.9a. This displays a handle $\{X, Y\}$ immediately followed by a cross-cap Z. Split c and d to create a red-yellow bigon W and let U_1 denote the resulting graph. (See Figure II.9b.) In U_1, a_1 and a_2 are in distinct red-yellow bigons, and so the operation of cancellation of a red-blue bigon is applicable to Z. We apply it, and let U_2 denote the resulting graph. (See Figure II.9c.) In U_2 split c' and d' to create a red-blue bigon V, and let U_3 denote the resulting gem. (See Figure II.9d.) In U_3, a_4 and a_6 are in distinct red-yellow bigons, and so the operation of cancellation of a red-blue bigon is applicable to X. We apply it, and let U_4 denote the resulting graph. (See Figure II.9e.)

At this stage we still have a unitary gem U_4. We have an assembled cross-cap V, an unassembled cross-cap Y and a cap W. The original assembled handles and cross-caps which have edges in P are preserved.

Our next step is to replace the cross-cap Y by an assembled cross-cap, as in Lemma II.6. The assembled cross-caps and handles of U_4 are clearly preserved. The red-blue bigon W is transformed into another cross-cap by Lemma II.7. Finally, we assemble this cross-cap, too. We thus obtain a unitary gem, still of the form in Lemma II.8, but with one handle fewer and two more cross-caps. Repetition of the above procedure leads us to a unitary gem in which all red-blue bigons are cross-caps.

Only if all red-blue bigons in U belong to assembled handles is the above procedure inapplicable. Hence the lemma follows. ❑

a)

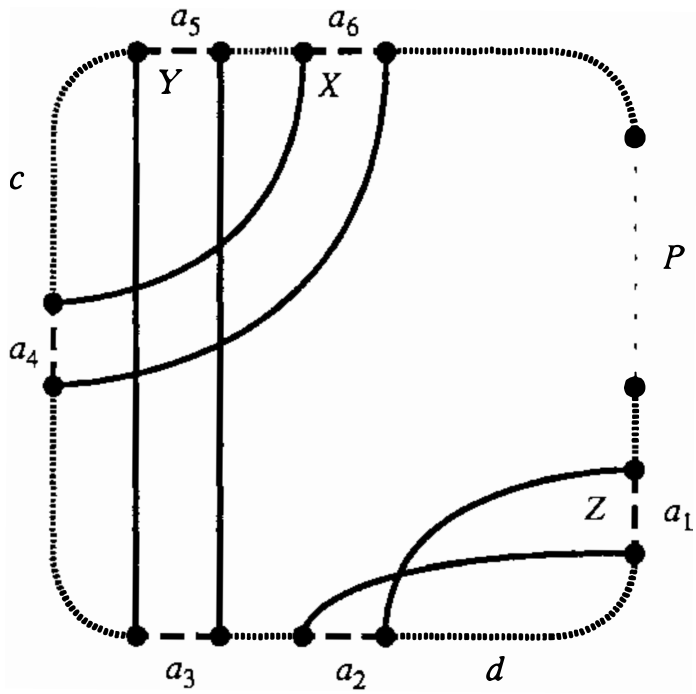

b)

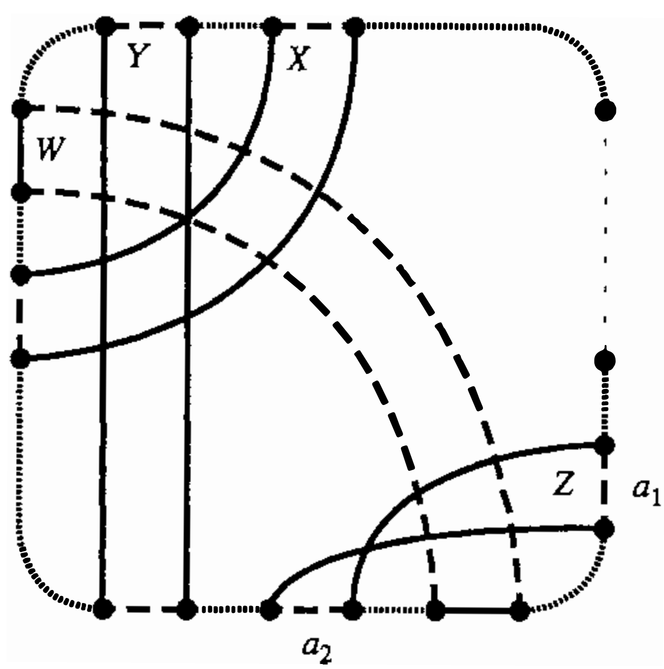

c)

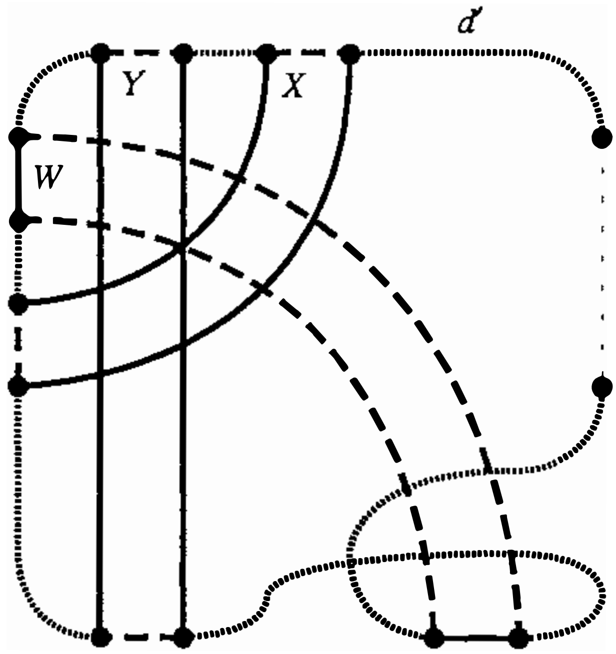

d)

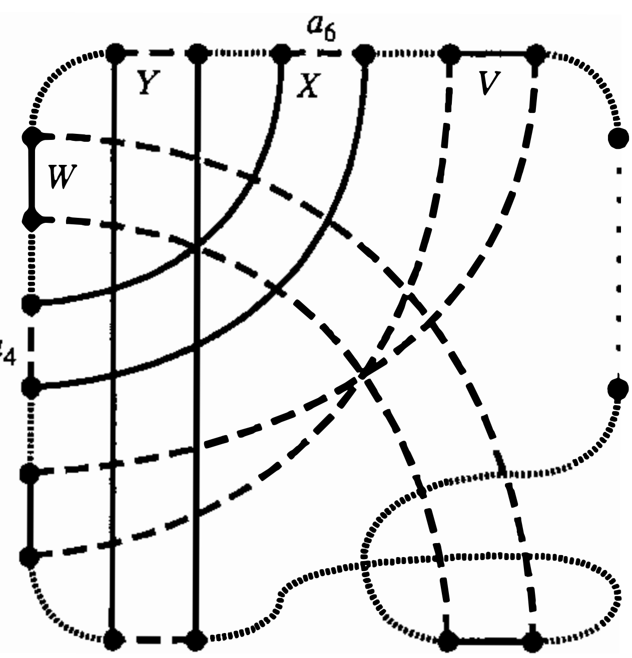

e)

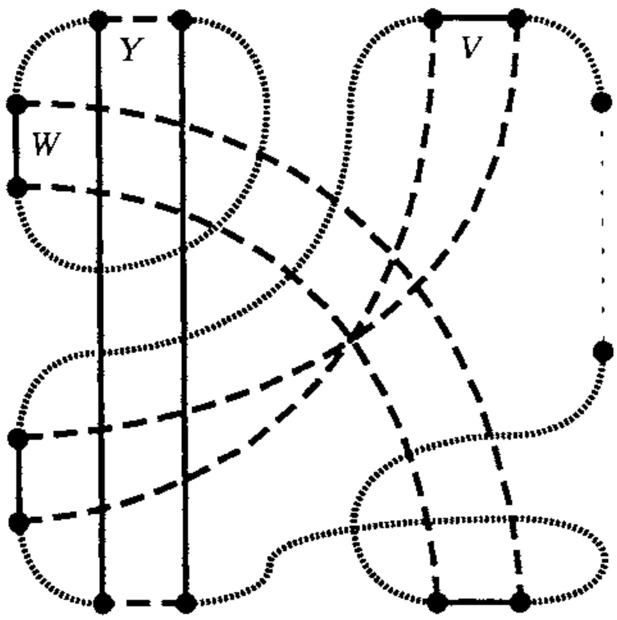

FIGURE II.9

5. CANONICAL GEMS

Let us define a *canonical* 3-graph as either the trivial 3-graph or a unitary gem U in which $Y(U)$ consists entirely of assembled cross-caps or members of assembled handles.

The trivial 3-graph is orientable and so is a unitary canonical 3-graph whose red-blue bigons are members of handles. The *genus* of such a 3-graph is the number of handles that it contains. The *genus* of the trivial 3-graph is defined to be zero. On the other hand, a unitary canonical 3-graph whose red-blue bigons are cross-caps is non-orientable. The *cross-cap number* of such a 3-graph is the number of cross-caps that it contains.

We observe that an orientable canonical 3-graph can be constructed with an arbitrary non-negative integer as genus, and that a non-orientable one can be constructed with an arbitrary positive integer as cross-cap number. There are no other possibilities. Hence we have the following lemma.

LEMMA II.10. *There is at most one orientable canonical* 3-*graph with given genus and one non-orientable canonical* 3-*graph with given cross-cap number.* ❑

LEMMA II.11. *An orientable canonical* 3-*graph of genus* g *has Euler characteristic* $2-2g$. *A non-orientable canonical* 3-*graph of cross-cap number* k *has Euler characteristic* $2-k$.

Proof. This follows from the fact that the number of red-blue bigons is $2g$ in the first case and k in the second. ❑

LEMMA II.12. *There is exactly one canonical* 3-*graph on each surface.*

Proof. By Lemmas II.5 and II.9 there exists a canonical 3-graph on each surface. Its uniqueness follows from Lemmas II.10 and II.11. ❑

LEMMA II.13. *Let* K *and* J *be* 3-*graphs on the same surface. Then* K *and* J *are* μ-*equivalent.*

Proof. Let U be the canonical 3-graph on the surface. Then each of K and J can be transformed into U by a sequence of μ-moves by Lemmas II.5 and II.9. The lemma follows. ❑

Theorem II.2 now follows from Lemmas II.3 and II.13. The genus of an orientable canonical gem is also called the *genus* of its surface, and of any other 3-graph on that surface. Likewise, the cross-cap number of an non-orientable canonical gem is also called the *crosscap number* of its surface, and of any other 3-graph on that surface. A surface is called *orientable* or *non-orientable* according as the 3-graphs on it are orientable or non-orientable. We can now say that there is just one orientable surface $\mathbb{S}_g$ whose genus is a given non-negative integer g, and just one non-orientable surface $\mathbb{N}_k$ whose cross-cap number is a given positive integer k, and moreover there are no other surfaces.

6. CONCLUSION

We have established the classification of surfaces by means of simple operations (dipole cancellations and creations) on 3-graphs. These operations enabled us to reduce any given 3-graph to a simple canonical form. This observation provides us with a possible approach to proving results about cubic graphs with a proper edge colouring in three colours: first prove the required result for the canonical forms, and then prove that the result is preserved under the dipole cancellation and creation operations.

We would like to mention that Vince (private communication) has a similar method for obtaining the classification of surfaces. It is of interest to note that his canonical 3-graphs are different from ours.

Chapter III

THE BOUNDARY AND FIRST HOMOLOGY SPACES OF A 3-GRAPH

1. INTRODUCTION

In recent years, several authors have investigated topological graph theory from a combinatorial viewpoint. In particular, graph theoretic versions of the Jordan curve theorem have been proved in [13, 24, 25, 31]. In some of these papers ([13, 31]) the development is in terms of 3-graphs. In the present chapter, this work is extended to a graph theoretic version of the theorem that the first Betti number of a surface is the largest number of closed curves that can be drawn on the surface without dividing it into two or more regions. Again the treatment is in terms of 3-graphs.

2. THE BOUNDARY SPACE

Some of the following concepts appeared originally in the work of Stahl [23, 24, 25] on permutation pairs.

Let K be a 3-graph. A non-empty set C of edges of K is called a *b-cycle* if C is the union of disjoint circuits with at least one blue edge in each. A set S of b-cycles *induces* a b-cycle C if each blue edge of C is an element of $\bigcup S$. In the special case where $S = \{D\}$ for a b-cycle D, we also say that D *induces* C. For example, any b-cycle induces itself. The *boundary space* of K is the subspace of $Z(K)$ spanned by the set of bigons of K. A b-cycle is said to *separate* if it induces a b-cycle which is a member of the boundary space. A set S of b-cycles is said to *separate* if it induces a b-cycle which separates. A b-cycle is *connected* if it is a circuit.

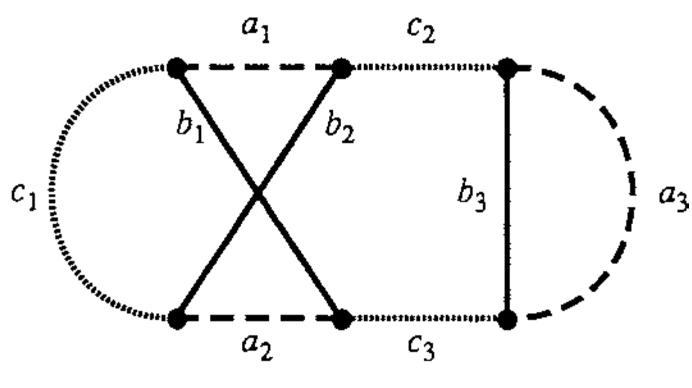

FIGURE III.1

EXAMPLE III.1. *Consider the 3-graph of Figure* III.1 *and let $C = \{b_1, c_3, b_3, c_2, a_1\}$. The b-cycle C separates since it induces the b-*

cycle $\{b_3, a_3\}$. *However* C *is not a member of the boundary space since it is not a sum of bigons.* ❑

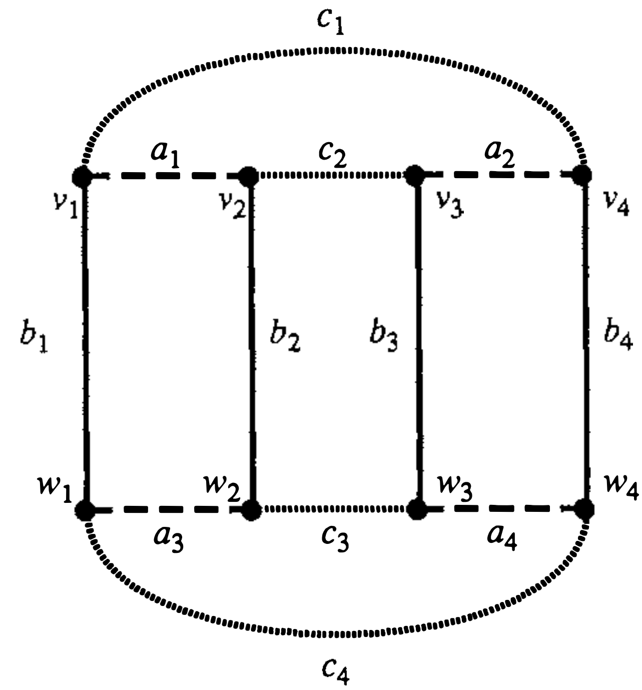

FIGURE III.2

3. SEMICYCLES

Let L be a subgraph or a set of edges of K. We write $\beta(L)$ and $\rho(L)$ for the set of blue and red edges respectively in L. Moreover, $N(L)$ denotes the set of red-yellow bigons that contain an edge in L. Hence $N(L) \subseteq B(K)$.

Let C be a b-cycle. We say that $N(C)$ is the *necklace* of C. The elements of its necklace are the *beads* of C. The *poles* of a bead B (with respect to C) are the vertices of B incident with a blue edge of

C. If C is connected and each bead has just two poles, then C is a *semicycle.*

EXAMPLE III.2. *Consider the* 3-*graph of Figure* III.2 *and let* C *be the connected b-cycle* $\{b_2, c_2, a_2, b_4, a_4, c_3\}$. *Let* $B_1 = \{a_1, c_2, a_2, c_1\}$ *and* $B_2 = \{a_3, c_3, a_4, c_4\}$. *Then* $\beta(C) = \{b_2, b_4\}$, $\rho(C) = \{a_2, a_4\}$ *and* $N(C) = \{B_1, B_2\}$. *The poles of* B_1 *with respect to* C *are* v_2 *and* v_4. *Likewise the poles of* B_2 *with respect to* C *are* w_2 *and* w_4. *Hence* C *is a semicycle. The connected b-cycle* $\{a_1, b_2, c_3, b_3, a_2, b_4, c_4, b_1\}$ *is not a semicycle since the poles of* B_1 *are* v_1, v_2, v_3 *and* v_4. ❑

Note that if a semicycle C induces a b-cycle C', then C' must be a semicycle with the same blue edges as C.

The concept of a semicycle was introduced by Stahl [24] in the setting of permutation pairs, but the motivation for it is best explained by considering the case where K is a gem. As we indicated earlier, K then corresponds to a 2-cell embedding, in a closed surface $\mathbb{S}$, of a graph G. Under this interpretation, the beads of a semicycle C of K correspond to vertices of G. The requirements that C should be connected and have a blue edge, and that each bead should have just two poles, reveal that C corresponds to a circuit D of G or a path of length 1. In the former case, observe that if D divides $\mathbb{S}$ into two regions, then D is the sum of the boundaries of the faces inside one of those regions. (We consider an edge to belong to the boundary of a face F if and only if it separates two distinct faces, one of which is F.) If R is either of those regions, then C is the sum of the blue-yellow bigons of K corresponding to the faces inside R, the red-yellow bigons corresponding to vertices in the interior of R, the red-blue bigons corresponding to edges in the interior of R, and possibly some of the red-yellow and red-blue bigons corresponding to the vertices and edges, respectively, of D. We infer that C is a member of the boundary space, and hence separates. Conversely, if the semicycle C separates, then C is a sum of bigons, and it follows that D divides $\mathbb{S}$ into two regions. The vertices, edges and faces in the interior of one of these regions correspond to those bigons in the sum which are not beads of C or red-blue bigons which meet C.

On the other hand, suppose that C corresponds to a path Q in G of length 1. Then C has just two beads. Either C is the red-blue bigon corresponding to the unique edge of Q, or C is the sum of this bigon and one or both of the beads. In any case, C separates.

Note also that each circuit of G has a non-empty family of semicycles of K which correspond to it.

4. B-INDEPENDENT SETS OF B-CYCLES

The members of a set S of b-cycles are *b-dependent* if there exists $C \in S$ induced by $S - \{C\}$. The b-cycles in S are *b-independent* if they are not b-dependent.

EXAMPLE III.3. *Consider the gem of Figure* III.3. *Let* $C_1 = \{b_5, c_2, a_4, c_3, b_2, c_4\}$ *and* $C_2 = \{b_4, c_6, b_1, c_1, a_3, c_2\}$. *Hence* $S = \{C_1, C_2\}$ *is a set of two b-independent semicycles since the blue edge* $b_4 \in C_2$ *is not in* C_1 *and similarly the blue edge* $b_5 \in C_1$ *is not in* C_2. *Now let* $C_3 = \{a_1, b_2, a_2, b_1\}$. *Then* $\{C_1, C_2, C_3\}$ *is a set of b-dependent semicycles since* C_3 *is induced by* $\{C_1, C_2\}$. ❑

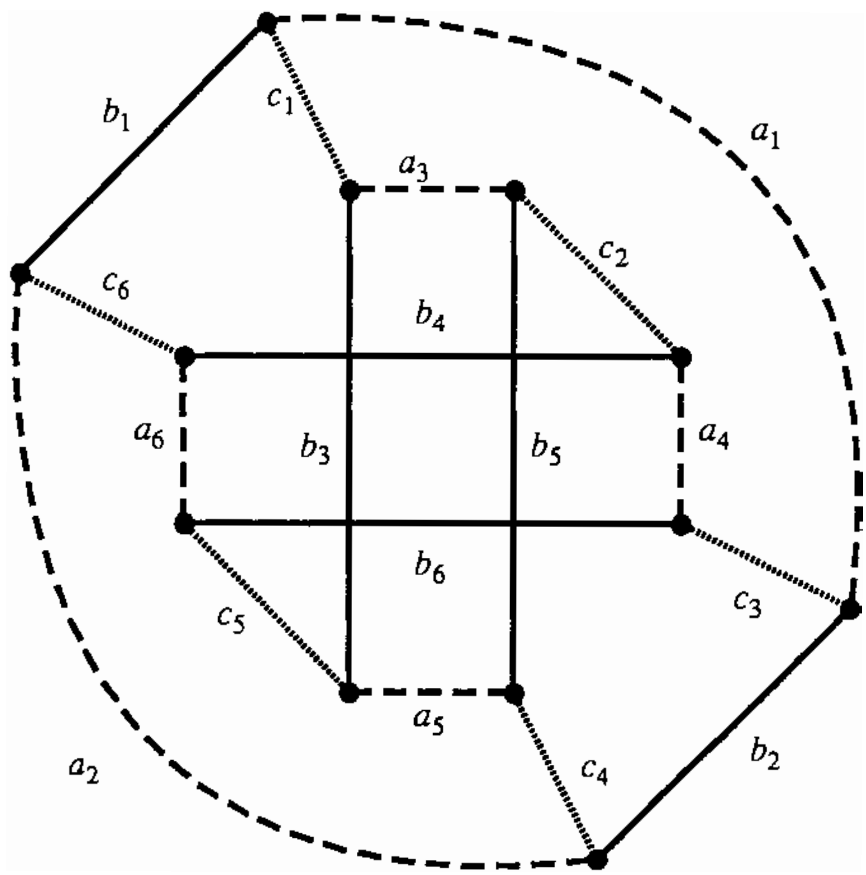

FIGURE III.3

One theorem in this chapter, Theorem III.15, asserts that in any 3-graph K the cardinality of a maximum set of b-independent semicycles which does not separate is the dimension of the *first* *homology space* $\boldsymbol{H}(K)$ of K, the orthogonal complement in $Z(K)$ of the boundary space of K. We henceforth use $h(K)$ to denote the dimension of $\boldsymbol{H}(K)$. In section 5, we explain the topological significance of this result and show how it can be used to deduce the version of the Jordan curve theorem which appeared in [24].

5. LINKING THE SIDES OF SEMICYCLES

The *conjugate* C^* of a b-cycle C is defined as $C + (\bigcup N(C))$, and is also a b-cycle. We define $I(C) = C \cap (\bigcup N(C))$ and $O(C) = C^* \cap (\bigcup N(C))$. We note that $\{I(C), O(C)\}$ is a partition of $\bigcup N(C)$. We call $I(C)$ and $O(C)$ the *sides* of C. This terminology is suggested by the interpretation of a gem as a model for an embedding of a graph.

EXAMPLE III.4. *Consider* C_1 *and* C_2 *of Example* III.3. *Then* $I(C_1) = \{c_2, a_4, c_3, c_4\}$ *and* $O(C_1) = \{a_3, c_1, a_1, a_2, c_6, a_6, c_5, a_5\}$. *Similarly,* $I(C_2) = \{c_6, c_1, a_3, c_2\}$ *and* $O(C_2) = \{a_6, c_5, a_5, c_4, a_2, a_1, c_3, a_4\}$. ❑

Note that C^* separates if and only if C does.

Now let P be a non-trivial path in K whose terminal edges are blue. For such paths we define necklaces, beads, poles and conjugates as for b-cycles, and we use analogous notation for these concepts. If each bead of the necklace $N(P)$ has just two poles, we call P a *semipath* of K. In addition, a red-yellow bigon is called a *terminal bead* of P if it contains a terminal vertex of P. P therefore has just one or two terminal beads. The elements of $N(P)$ are sometimes called *internal beads* of P. Each internal bead of a semipath has just two poles, and so none of them is terminal.

EXAMPLE III.5. *Consider the gem of Figure* III.3 *and let* $P = \{b_1, c_6, b_4\}$. *Then P is a semipath with one terminal bead* $\{c_1, a_3, c_2, a_4, c_3, a_1\}$ *and one internal bead* $\{c_6, a_6, c_5, a_5, c_4, a_2\}$. ❑

If K is a gem representing an embedding of a graph G, then a semipath P represents a path or circuit in G according to whether P has two terminal beads or just one. If P represents a path, then the internal and terminal beads of P correspond to the internal and terminal vertices, respectively, of the path. If P represents a circuit, then the internal beads and the unique terminal bead correspond to the vertices of the circuit.

A b-cycle C and a semipath P are said to *miss* if they are disjoint and no internal bead of P is a bead of C. (If K is a gem, C a semicycle in K corresponding to a circuit D in $G(K)$, and P a semipath in K corresponding to a path Q in $G(K)$, then C and P miss if and only if D and Q are disjoint and no internal vertex of Q is a vertex of D. If Q is a circuit rather than a path, then C and P miss if and only if D and Q are disjoint and no vertex of Q corresponding to an internal bead of P is a vertex of D.) If C and P miss, then P is said to *link* the sides of C if one terminal vertex is in $VI(C)$ and the other is in $VO(C)$. (We note that the terminal vertices of P cannot be poles of a bead of C, since $P \cap C = \varnothing$.)

EXAMPLE III.6. *The semipath P of Example* III.5 *links the sides of the semicycle C_1 of Example* III.3. ❑

A b-cycle C is *normal* if no blue edge of C is adjacent to edges of C with distinct colours. Clearly the conjugate of a normal b-cycle is also normal.

Theorem III.8, which is proved in Section 6, asserts that a b-cycle separates if and only if it induces a normal b-cycle whose sides are not linked by a semipath. This theorem also has a topological interpretation, which may be discerned from the following considerations. Let K be a gem corresponding to a 2-cell imbedding of a graph G in a closed surface $\mathbb{S}$, and let D be a circuit in G corresponding to a normal semicycle C in K. Then the two sides of C determine a pair of complementary subsets E_1 and E_2 of the set $\partial_G VD$. Each of these subsets consists of all the edges in $\partial_G VD$ whose corresponding red-blue bigons in K meet a given side of C. Intuitively, E_1 and E_2 may be thought of as representing the sides of D in the embedding of G, since C is normal. Now let P be a semipath in K of length greater than 1. If P corresponds to a path Q in G, then P links the sides of C in K if and only if Q joins two vertices of D, meets both E_1 and E_2 and none of its internal vertices is a vertex of D. On the other hand, if Q is a circuit rather than a path then P links the sides of C if and only if Q meets both E_1 and E_2 and has just one vertex in common with D.

Intuitively, Theorem III.8 therefore asserts that C separates if and only if D divides $\mathbb{S}$ into two regions in such a way that if Q is a path in G with no internal vertex in VD, or a circuit in G having at most one vertex in VD, then the vertices and edges of Q that are not in D are collectively confined to one region.

EXAMPLE III.7. *Consider the b-cycle C of Example* III.1. *C induces the normal b-cycle $C' = \{b_3, a_3\}$. Moreover $I(C') = \{a_3\}$ and $O(C') = \{c_3, a_2, c_1, a_1, c_2\}$ and therefore it is evident that no semipath links the sides of C'. This result of course agrees with Theorem* III.8. ❑

If C is a semicycle of K which induces another b-cycle C', then C' is a semicycle such that $\beta(C') = \beta(C)$. Now suppose that K is a gem corresponding to a 2-cell embedding of a graph G in a closed surface $\mathbb{S}$. It follows from the above observation that C and C' correspond to the same path or circuit of G. Note also that a set of semicycles in a gem is b-independent if and only if the corresponding paths and circuits in G have the property that each contains an edge not in the union of the others. Thus a necessary and sufficient condition for a set of semicycles to be a set S of b-independent semicycles which does not separate is for S to correspond to a set of circuits in G which are collectively drawn so as not to divide $\mathbb{S}$ into two or more regions and have the property that each contains an edge not in any of the others. According to Theorem III.15, the cardinality of a maximum set of such semicycles is the dimension of the first homology space of K. It is shown in [13] that the dimension of the boundary space is $r(K) - c(K)$. Therefore the dimension $h(K)$ of the first homology space $\boldsymbol{H}(K)$ is

$$\begin{aligned} |EK| - |VK| + c(K) - (r(K) - c(K)) &= 2c(K) - r(K) + \frac{|VK|}{2} \\ &= 2c(G) - \chi(\mathbb{S}) \end{aligned}$$

since K is cubic and $c(K) = c(G)$. This number is the first Betti number of $\mathbb{S}$. These observations show that Theorem III.15 is a graph theoretic version of the theorem that the first Betti number of a surface is the largest number of closed curves that can be drawn on the surface without dividing it into two or more regions.

The considerations above also show that $\mathbb{S}$ is the sphere if and only if the dimension of the first homology space of K is 0. Thus G is planar if and only if the set of bigons of K spans $\boldsymbol{Z}(K)$.

In order to obtain a graph theoretic version of the Jordan curve theorem, let us suppose first that $\mathbb{S}$ is the sphere. Since the sphere is orientable, K is bipartite. Let C be a normal semicycle in K. By Theorem III.15, $\{C\}$ separates, and therefore induces a b-cycle which separates. This b-cycle induces a normal b-cycle D whose sides are not linked by a semipath. D therefore separates. But D, being normal and induced by $\{C\}$, must be C or C^*. In either case we conclude that C separates.

On the other hand, suppose that K is bipartite but that $\mathbb{S}$ is not the sphere. Then K has a semicycle C which does not separate. C therefore does not induce a b-cycle which is a sum of bigons. But it is easy to show that C induces a normal semicycle. We begin by observing that if b is a blue edge of C adjacent to edges of C with distinct colours, then by adding to C a bead containing one of these edges we obtain a semicycle induced by C in which the edges adjacent to b are of the same colour. By applying this construction to all but one of the blue edges of C, we can find a semicycle C' which is induced by C and has the property that at most one blue edge of C' is adjacent to edges of C' with distinct colours. The requirement that $|C'|$ be even, because K is bipartite, then forces C' to be normal. Since C' is induced by C, it cannot be a sum of bigons. Any b-cycle induced by C' is obtained from C' by the addition of beads of C', and therefore is not a sum of bigons either. We conclude that C' does not separate.

In summary, we have shown that if K is bipartite then G is planar if and only if every normal semicycle of K separates. This is the graph theoretic version of the Jordan curve theorem which appeared in [13]. It implies another version, due to Stahl [24].

In the case where K is non-bipartite, it may not be possible to construct a normal semicycle, separating or non-separating. Consider the gem in Figure III.4. The four possible semicycles in this gem are $\{b_1, c_1, a_2\}$ and $\{b_2, c_2, a_2\}$ and their conjugates. In all cases the semicycle is not normal.

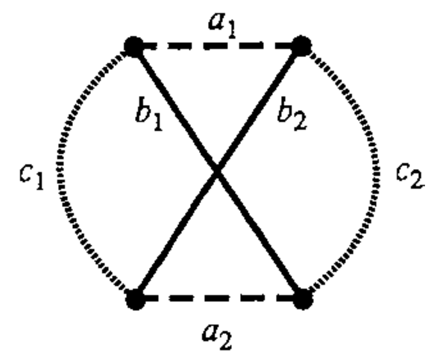

FIGURE III.4

6. A CONDITION FOR A B-CYCLE TO SEPARATE

THEOREM III.8. *A necessary and sufficient condition for a b-cycle D to separate a 3-graph K is for D to induce a normal b-cycle D′ such that no semipath links the sides of D′.*

The proof of this theorem consists of several lemmas. We may assume that K is connected, since the general case follows by applying the theorem to each component separately.

LEMMA III.9. *Let $C = \sum U$ where U is a set of red-blue and blue-yellow bigons in a 3-graph K. Let b be a blue edge, not in $\beta(C)$, joining vertices v and w. Then either $\{v, w\} \subseteq VC$ or $\{v, w\} \cap VC = \varnothing$.*

Proof. Suppose $v \in VC$. Let a_1 and c_1 be the red and yellow edges respectively incident on v. Let a_2 and c_2 be the red and yellow edges respectively incident on w. Since $a_1 \in C$ we have $b \in \bigcup U$. Hence $a_2, c_2 \in C$ and therefore $w \in VC$. Similarly one can show that $w \notin VC$ when $v \notin VC$. ❑

LEMMA III.10. *If a b-cycle D separates then D induces a normal b-cycle D' such that no semipath links the sides of D'.*

Proof. Since D separates, D induces a b-cycle $\sum \boldsymbol{U}$ where $\boldsymbol{U}$ is a set of bigons. Let $\boldsymbol{U}_1$ and $\boldsymbol{U}_2$ denote the set of red-blue and blue-yellow bigons, respectively, included in $\boldsymbol{U}$, and let $C = \sum(\boldsymbol{U}_1 \cup \boldsymbol{U}_2)$. Now consider the cycle $D' = C + \bigcup \boldsymbol{B}$ where $\boldsymbol{B}$ is the set of red-yellow bigons included in C. (We include $\bigcup \boldsymbol{B}$ in the above sum of bigons to ensure that all circuits in D' contain a blue edge.) Since $\varnothing \neq \beta(\sum \boldsymbol{U}) = \beta(D')$, then D' is a b-cycle induced by D. Also, the two edges of D' adjacent to a given blue edge of D' must belong to the same bigon, and hence D' is a normal b-cycle.

We claim that no semipath links the sides of D'. Assume by way of contradiction that P is a semipath that links the sides of D'. Let v denote the terminal vertex of P that is in $VI(D')$, and let b be the blue edge incident on v (b is therefore in P). Let b join v to w and let B denote the red-yellow bigon that contains w. By Lemma III.9, $w \in VC$ and therefore $B \in N(C)$. If $B \in N(D')$, then $P = \{b\}$ and $v, w \in VI(D')$, a contradiction. Hence we conclude that $B \in \boldsymbol{B}$ and $B \subseteq C$. Therefore the blue edge of $P - \{b\}$ incident on a vertex of VB must have both end vertices in VC. Proceeding inductively along P, we find that the terminal vertex x of P other than v must lie in VC. Since the red-yellow bigon B' that contains x must be in $N(D')$ it follows that $x \in VI(D')$, a contradiction. The lemma follows. ❑

Lemma III.10 proves half of Theorem III.8. Accordingly we henceforth assume that D is a b-cycle such that no b-cycle induced by D is a sum of bigons. Let D' be an arbitrary normal b-cycle induced by D. We shall show that the sides of D' are linked by a semipath.

Let $K^{\dagger}$ be a graph whose vertices are the red-blue and blue-yellow bigons of K and whose edges are the blue edges of K. Any edge $b \in EK^{\dagger}$ is to join the two bigons containing b in K. Clearly $K^{\dagger}$ is connected since K is connected.

LEMMA III.11. *The graph $K^{\dagger} - \beta(D')$ is connected.*

Proof. Suppose $K^{\dagger} - \beta(D')$ is unconnected, and let $L^{\dagger}$ be a component of $K^{\dagger} - \beta(D')$. Let us consider the b-cycle $D'' = \sum VL^{\dagger} + \bigcup B$ where B is the set of red-yellow bigons included in $\sum VL^{\dagger}$. (Recall that $\sum VL^{\dagger}$ is the symmetric difference of the bigons that constitute the vertex set of $L^{\dagger}$.) By the construction $\beta(D'') \subseteq \beta(D') \subsetneq \beta(D)$ and therefore D'' is a b-cycle induced by D that is a sum of bigons. This contradicts our assumption. The lemma follows. ❑

If there is an edge b of $\beta(K) - \beta(D')$ with an end vertex in $VI(D')$ and one in $VO(D')$ then $\{b\}$ links the sides of D'. Henceforth we suppose there is no blue edge with this property, and partition the set $\beta(K) - \beta(D')$ into three classes: the set I of edges with an end vertex in $VI(D')$, the set O of edges with an end vertex in $VO(D')$, and the set M of edges incident with no vertex in $VI(D') \cup VO(D')$.

LEMMA III.12. *The edge sets I and O are non-empty.*

Proof. Assume that I is empty. Then there is no vertex of VD' incident on a red edge and a yellow edge of D'. Since D' is normal it must therefore be the sum of a disjoint set of red-blue and blue-yellow bigons, a contradiction. A similar argument applied to the conjugate of D' shows that $O \neq \emptyset$. ❑

LEMMA III.13. *Let P be a path with terminal vertices v and w and blue terminal edges. Suppose that the red-yellow bigons containing v and w are not in N(P). Then there exists a semipath P′, joining v and w, such that $\beta(P') \subseteq \beta(P)$ and $N(P') \subseteq N(P)$.*

Proof. We use induction on $|\beta(P)|$. If $|\beta(P)| = 1$ then P is the required semipath. Now suppose the lemma holds for all paths with fewer than $|\beta(P)|$ blue edges, where $|\beta(P)| > 1$. Let b denote the blue terminal edge of P incident on v. Let B denote the red-yellow bigon in $N(P)$ containing a vertex x incident with b. (B exists since

$|\beta(P)| > 1$.) Let y be the vertex of $VP \cap VB$ that minimises $|P[w, y]|$. Thus $Q = P[w, y]$ is a path with fewer blue edges than P. Furthermore $B \notin N(Q)$ by the choice of y. By the inductive hypothesis, there exists a semipath P' with terminal vertices y and w such that $\beta(P') \subseteq \beta(Q) \subseteq \beta(P)$ and $N(P') \subseteq N(Q) \subseteq N(P)$. Let Q' be a path included in B which joins x and y. Then $P' \cup Q' \cup \{b\}$ is the required semipath. ❑

LEMMA III.14. *There exists a semipath that links the sides of D'.*

Proof. Case i) Suppose there exists a vertex Y in $VK^{\dagger}$ incident on an edge $b \in O$ and an edge $b' \in I$. We may choose b and b' so that they are terminal edges of a path P included in the bigon Y such that $\beta(P) - \{b, b'\} \subseteq M \cup D'$. Let $\beta(P) = \{b_1, b_2, ..., b_n\}$, where $b_1 = b$ and $b_j \in P[b_i, b_k]$ whenever $i < j < k$. Thus $b_n = b'$. For each positive integer $i < n$ let a_i be the edge of P joining an end-vertex w_i of b_i to an end-vertex v_{i+1} of b_{i+1}, and let d_i and c_{i+1} be the edges of $EK - Y$ incident on w_i and v_{i+1} respectively. (See Figure III.5.)

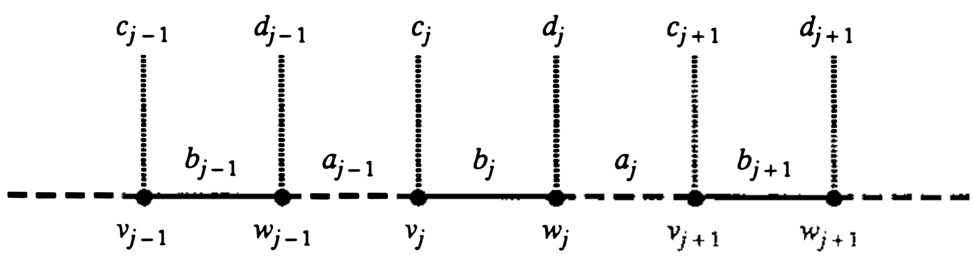

FIGURE III.5

Suppose $b_j \in D'$ for some $j < n$. Choose j to be as small as possible subject to this requirement. We have $j > 1$ since $b_1 \notin D'$. Hence $b_{j-1} \in M \cup O$. It follows that $w_{j-1} \notin VD'$, and so $c_j \in D'$. Therefore $d_j \in D'$ since D' is normal. If $j + 1 < n$ then $b_{j+1} \notin M$; hence $\{c_{j+1}, b_{j+1}, d_{j+1}\} \subseteq D'$. By induction $\{c_{n-1}, b_{n-1}, d_{n-1}\} \subseteq D'$. Therefore $a_{n-1} \notin D'$. Since $b' \notin D'$, we obtain the contradiction that $b' \in O$. Hence $\{b_2, b_3, ..., b_{n-1}\} \subseteq M$, and so $N(P) \cap N(D') = \varnothing$. By Lemma III.13, there exists a semipath P' with terminal edges b_1 and b_n such that P' misses D'. Therefore P' links the sides of D', as required.

Case ii) Suppose there is no bigon in $VK^\dagger$ incident on an edge in I and on an edge in O. By Lemma III.11 and Lemma III.12 there exists a path $P^\dagger \subseteq M$ in $K^\dagger - \beta(D')$ with terminal vertices Y_1 and Y_2 such that $Y_1 \cap O \neq \varnothing$ and $Y_2 \cap I \neq \varnothing$. We may assume $P^\dagger$ chosen so that $\beta(L) \cap (I \cup O) = \varnothing$, where $L = \bigcup(VP^\dagger - \{Y_1, Y_2\})$. (Recall that $VP^\dagger$ is the set of bigons in K that make up the vertices in $P^\dagger$.) This choice guarantees that $N(Z) \cap N(D') = \varnothing$ for each internal vertex Z of $P^\dagger$. Let P_1 be a path in Y_1 with blue terminal edges $b \in O$ and $b_1 \in (L \cup Y_2) \cap Y_1$ such that the blue internal edges of P_1 are edges in M. Clearly $N(P_1) \cap N(D') = \varnothing$, for otherwise an internal blue edge of P_1 would not be in M. Similarly let P_2 be a path in Y_2 with blue terminal edges $b' \in I$ and $b_2 \in (L \cup Y_1) \cap Y_2$ such that the blue internal edges of P_2 are edges in M. Again, $N(P_2) \cap N(D') = \varnothing$. It follows that there is a path P in K, with terminal edges b and b', such that $N(P) \cap N(D') = \varnothing$. By Lemma III.13 we can construct a semipath P' with terminal edges b and b' that misses D'. P' links the sides of D', as required. ❑

7. FUNDAMENTAL SETS OF SEMICYCLES

A set of m b-independent semicycles that does not separate is said to be an *m-fundamental* set. In this section we shall show that the size of a maximum m-fundamental set is

$$h(K) \;=\; 2c(K) - r(K) + \frac{|VK|}{2}.$$

THEOREM III.15. *If K is a 3-graph then the maximum size of an m-fundamental set is $h(K)$.*

We devote this section to a proof of Theorem III.15. If $K_1, K_2, ..., K_{c(K)}$ denote the components of K then

$$\sum_{i=1}^{c(K)} h(K_i) = \sum_{i=1}^{c(K)} \left(2 - r(K_i) + \frac{|VK_i|}{2} \right)$$

$$= h(K).$$

Therefore if Theorem III.15 holds for each K_i then it holds for K. We henceforth assume K to be connected.

LEMMA III.16. *If S is an m-fundamental set then $m \leq h(K)$.*

Proof. Let C be a semicycle in S and suppose a is a red or yellow edge contained in C. Let B denote the red-yellow bigon that contains a. Clearly $C' = C + B$ is a semicycle such that $a \notin C'$ and $\beta(C') = \beta(C)$. Let $S' = (S - \{C\}) \cup \{C'\}$. Since C is not induced by $S - \{C\}$ then C' is not induced by $S' - \{C'\}$, and therefore S' is a b-independent set. Furthermore $\beta(\bigcup S) = \beta(\bigcup S')$ and therefore S' is an m-fundamental set such that the number of semicycles in S' containing a is one less than in S. Proceeding inductively we obtain an m-fundamental set S'' such there exists an edge in each red-yellow bigon of K that is not in $\bigcup S''$.

Since S'' does not separate, each semicycle in S'' is a cycle that does not lie in the boundary space of K. Also, S'' is a linearly independent set of cycles since it is b-independent. Let T be a set of bigons in K comprising all the bigons except for exactly one arbitrarily chosen bigon. It is shown in [13] that T is a basis for the boundary space of K. If $m > h(K)$ then $T \cup S''$ is linearly dependent since $|T \cup S''| > \dim \boldsymbol{Z}(K)$ and hence there exists a cycle D belonging to the boundary space which is a sum of semicycles in S''. By the construction of S'', it is impossible for D to include a red-yellow bigon. We conclude that D must be a b-cycle. Moreover, D is induced by S'', which is a contradiction. Hence we conclude that $m \leq h(K)$. ❑

8. IMPLIED SEMICYCLES

The next two sections are concerned with the construction of a $(h(K))$-fundamental set S. This construction together with Lemma

III.16 gives us our theorem. It also implies that S is a basis for the first homology space and that $T \cup S$ is a basis for the cycle space of K, where T is a set of bigons in K comprising all the bigons except for exactly one arbitrarily chosen bigon.

Let K be a 3-graph with a blue 1-dipole b. Let C' be a semicycle in the 3-graph $K' = K - [b]$. The following uses the notation in Figure II.2. If C' does not meet A' then all the edges of C' are in K and we define $C = C'$. If C' meets A' then let x and y be the two poles of A' with respect to C'. Let $P' = A - \{a_1, c_1\}$ and $Q' = B - \{a_2, c_2\}$. Assume without loss of generality that $x \in VP'$ and consider the following two cases.

Case a) $y \in VP'$. Let P be a path in A that links x to y. Then we define $C = (C' - A') \cup P$.

Case b) $y \in VQ'$. Let P be a path in A that links x to v and Q be a path in B that links w to y. We define $C = (C' - A') \cup P \cup Q \cup \{b\}$.

In the cases presented above, C is clearly a semicycle in K such that $\beta(C) - \{b\} = \beta(C')$. We say that C is a semicycle (in K) *implied* by the semicycle C' (in K').

More generally, let $S' = \{C_1', C_2', ..., C_n'\}$ be a set of semicycles in K'. For each i let C_i be a semicycle in K implied by C_i', and let $S = \{C_1, C_2, ..., C_n\}$. We say that S is a set of semicycles (in K) *implied* by the set S' of semicycles (in K'). Clearly $\beta(\bigcup S) - \{b\} = \beta(\bigcup S')$.

LEMMA III.17. *Let K be a 3-graph with blue 1-dipole b, and let $K' = K - [b]$. If S' is an m-fundamental set in K' then a set S of semicycles in K implied by S' is an m-fundamental set in K.*

Proof. First we show that S is b-independent. Suppose not, and let C be a semicycle in S that is induced by $S - \{C\}$. Therefore $\beta(C) \subseteq \beta(\bigcup(S - \{C\}))$. Let C' denote the semicycle in S' that implies C. The fact that $\beta(C') = \beta(C) - \{b\} \subseteq \beta(\bigcup(S - \{C\})) - \{b\}$ $= \beta(\bigcup(S' - \{C'\}))$ implies that S' is not b-independent, a contradiction. Hence we conclude that S is b-independent.

The following uses the notation of Figure II.2. Let Y denote the red-blue bigon in K that includes $\{a_1, b, a_2\}$ and let Y' be the red-blue bigon $(Y - \{a_1, b, a_2\}) \cup \{a\}$ in K'. Similarly, let R denote the blue-yellow bigon in K that includes $\{c_1, b, c_2\}$ and let R' be the blue-yellow bigon $(R - \{c_1, b, c_2\}) \cup \{c\}$ in K'.

Suppose that S induces a b-cycle $D = \sum \boldsymbol{U}$ where $\boldsymbol{U}$ is a set of bigons. Let $\boldsymbol{U}_1$ and $\boldsymbol{U}_2$ denote the set of red-blue and blue-yellow bigons, respectively, included in $\boldsymbol{U}$. Let $D_1 = \sum(\boldsymbol{U}_1 \cup \boldsymbol{U}_2) + \bigcup \boldsymbol{B}$ where $\boldsymbol{B}$ is the set of red-yellow bigons included in $\sum(\boldsymbol{U}_1 \cup \boldsymbol{U}_2)$. Then D_1 is also a b-cycle induced by S that separates. If $Y \in \boldsymbol{U}_1$, then let $\boldsymbol{U}_1' = (\boldsymbol{U}_1 - \{Y\}) \cup \{Y'\}$; otherwise let $\boldsymbol{U}_1' = \boldsymbol{U}_1$. If $R \in \boldsymbol{U}_2$, then let $\boldsymbol{U}_2' = (\boldsymbol{U}_2 - \{R\}) \cup \{R'\}$; otherwise let $\boldsymbol{U}_2' = \boldsymbol{U}_2$. Let $D' = \sum(\boldsymbol{U}_1' \cup \boldsymbol{U}_2') + \bigcup \boldsymbol{B}'$ where $\boldsymbol{B}'$ is the set of red-yellow bigons included in $\sum(\boldsymbol{U}_1' \cup \boldsymbol{U}_2')$. If $D' = \varnothing$ then D would have exactly one blue edge, namely b. By the definition of a b-cycle, D

would consist of one circuit, comprising b and some red and yellow edges that belong to a red-yellow bigon. This contradicts the fact that b is a blue 1-dipole. Hence we conclude that D' is a b-cycle which separates K'. Moreover D' is induced by S' since $\beta(D') = \beta(D_1) - \{b\}$. However, this is a contradiction since S' does not separate. Hence we conclude that S is an m-fundamental set in K. ❑

Now suppose that K' is obtained from K by a finite sequence of blue 1-dipole cancellations, and that C' is a semicycle in K'. Then we apply the definition of an implied semicycle inductively to obtain a semicycle C in K that is *implied* by C'. Similarly, we speak of a set of semicycles in K' *implying* a set of semicycles in K. By Lemma III.17, if S' is an m-fundamental set in K' then the set S of semicycles in K implied by S' is an m-fundamental set in K.

9. 3-GRAPHS WITH JUST ONE RED-YELLOW BIGON

Suppose K to be a 3-graph with just one red-yellow bigon B. Let b be a blue edge in K joining vertices x and y. Let P be a path in B with terminal vertices x and y. The blue edge b can be used to define a semicycle in K, namely $C = \{b\} \cup P$. We say that C is a semicycle *formed* from b. Let T be a spanning tree of $K^{\dagger}$, and let $T' = EK^{\dagger} - ET$.

LEMMA III.18. $|T'| = h(K)$.

Proof. Since K has exactly one red-yellow bigon, we have $|VK^{\dagger}| = r(K) - 1$. The number of edges in a spanning tree for the graph $K^{\dagger}$ is $|VK^{\dagger}| - 1 = r(K) - 2$. Observe that

$$\begin{aligned} |EK^{\dagger}| &= |\beta(K)| \\ &= \frac{|VK|}{2}. \end{aligned}$$

Hence

$$\begin{aligned} |T'| &= |EK^{\dagger}| - |ET| \\ &= \frac{|VK|}{2} - (r(K) - 2). \ \square \end{aligned}$$

LEMMA III.19. *If K is a connected 3-graph with one red-yellow bigon then there exists a $(h(K))$-fundamental set in K.*

Proof. Let $j = h(K) = |T'|$, where T' is defined as above. For each $b_i \in T'$ where $i = 1, ..., j$, let C_i be a semicycle formed from b_i. Let $S = \{C_1, C_2, ..., C_j\}$. Hence $\beta(\bigcup S) = T'$. Since all the b_i's are distinct, S is a set of j b-independent semicycles in K. Assume that S induces a b-cycle which separates. Then S induces a b-cycle D of the form $\sum U$, for some set U of bigons in K. Since we may add the red-yellow bigon to U and still have a b-cycle induced by S which separates, we may assume that U does not contain the red-yellow bigon. Therefore, $\beta(D) = \partial_{K^{\dagger}}U$. Hence there must exist an edge in T that is in $\beta(D)$. However, this is impossible since $\beta(D) \subseteq \beta(\bigcup S) = T'$. Therefore S is a $h(K)$-fundamental set in K. $\square$

THEOREM III.15. *If K is a* 3-*graph then the maximum size of an m-fundamental set is* $h(K) = 2c(K) - \chi(K)$.

Proof. By Lemma III.16 we have $m \leq h(K)$. Hence we are required to show that there exists a $h(K)$-fundamental set in K. Cancel blue 1-dipoles from K one at a time until none is left, and let K' denote the resulting graph. Therefore K' has exactly one red-yellow bigon. By Lemma III.19 there exists a $h(K')$-fundamental set S' in K'. Since K' is obtained from K by a finite sequence of blue 1-dipole cancellations, by Lemma III.17 the set of semicycles in K implied by S' is a $h(K)$-fundamental set in K. ❑

EXAMPLE III.20. *We now illustrate the procedure implied by Lemmas* III.18 *and* III.19, *and Theorem* III.15, *with the* 3-*graph* L_1 *in Figure* III.6*a. Clearly* $L_2 = L_1 - [b_5]$, $L_3 = L_2 - [b_6]$, *and* $L_4 = L_3 - [b_4]$. *(See Figures* III.6*b, c and d.) We note that* L_4 *has just one bigon of each type, and hence the graph in Figure* III.6*e is* $L_4{}^{\dagger}$. *A suitable tree in* $L_4{}^{\dagger}$ *is* $\{b_1\}$. *Therefore, semicycles in* L_4 *formed from* b_2 *and* b_3 *do not separate.*

Let $C_1 = \{b_3, c_9, a_3, c_2\}$ *and* $C_2 = \{b_2, c_1, a_2, c_9\}$. *From, the proof of Theorem* III.19, $S = \{C_1, C_2\}$ *does not separate. A semicycle in* L_1 *implied by* C_1 *is* $C_3 = \{b_3, c_3, b_5, c_5, b_4, c_6, b_6, a_4\}$. *Similarly, a semicycle in* L_1 *implied by* C_2 *is* $C_4 = \{b_2, a_1, b_5, c_5, b_4, c_6, b_6, c_4\}$. $S' = \{C_3, C_4\}$ *does not separate* L_1. S' *is the largest* *possible set with this property since $\chi(L_1) = 0$. L_1 is a graph on the combinatorial torus.* ❑

a)

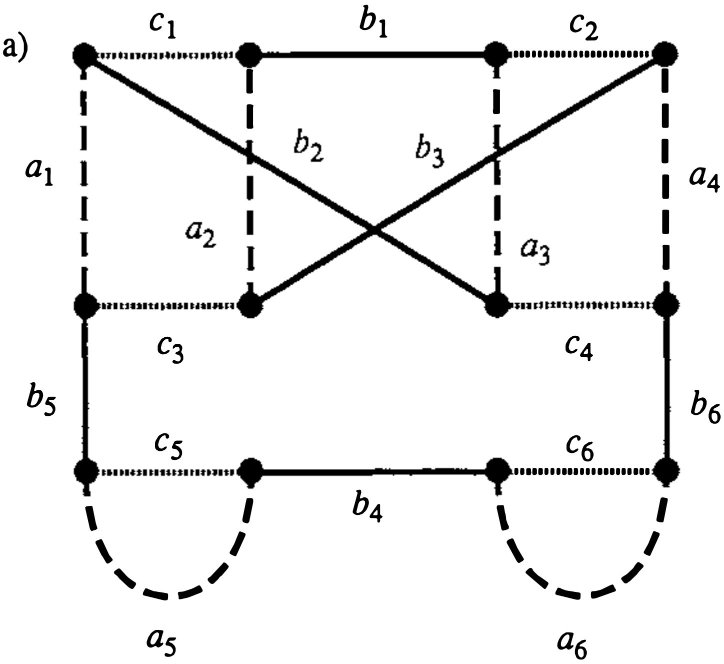

b)

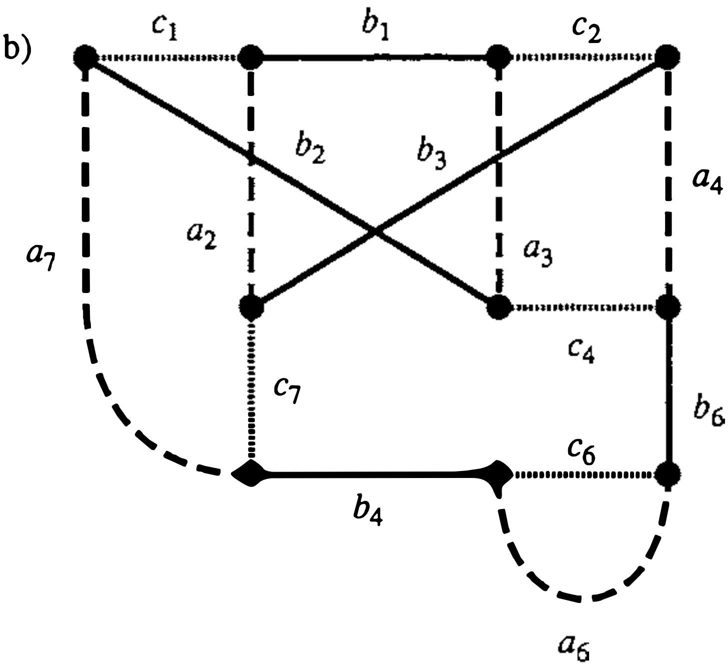

c)

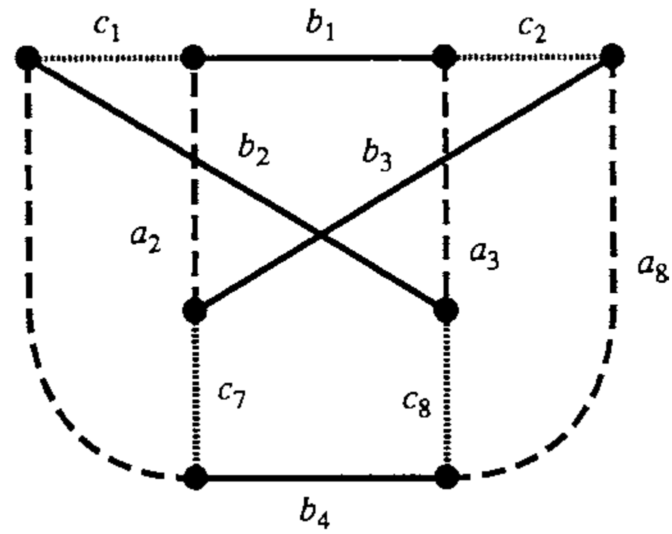

d)

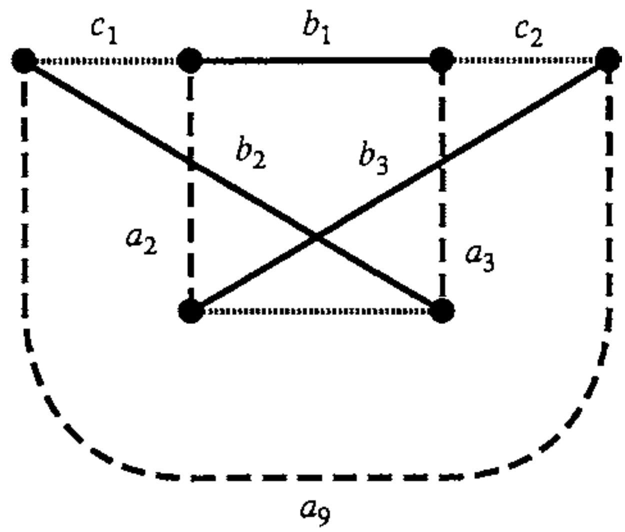

e)

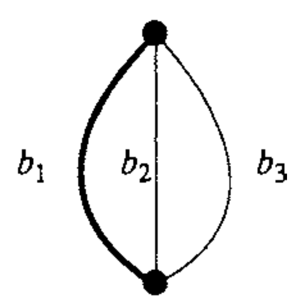

FIGURE III.6

Chapter IV

THE IMBEDDING DISTRIBUTION OF A 3-GRAPH

1. INTRODUCTION

A graph G may underlie many gems. This observation motivates the following definition. Let K and L be two 3-graphs. Suppose there exist bijections θ, φ, σ between $B(L)$ and $B(K)$, $\beta(L)$ and $\beta(K)$, and $\rho(L)$ and $\rho(K)$ respectively. Furthermore, suppose that

i) for any red-yellow bigon B in L and any red edge $a \in B$ we have $\sigma(a) \in \theta(B)$, and

ii) for any blue edge b adjacent to a red edge a we have $\varphi(b)$ adjacent to $\sigma(a)$.

Then K and L are *congruent*. Thus two gems are congruent if and only if a graph underlies them both. Moreover, if K and L are congruent then by condition ii) we have a bijection between the red-blue bigons of K and the red-blue bigons of L. Evidently, congruence is an equivalence relation.

EXAMPLE IV.1. *Consider the gems L and K in Figure* IV.1. *Let* B_1, B_2, B_3 *and* B_4 *denote the red-yellow bigons* $\{a_1, c_1, a_2, c_2\}$, $\{a_3, c_3, a_4, c_4\}$, $\{a_5, c_5, a_6, c_6\}$, *and* $\{a_7, c_7, a_8, c_8\}$ *respectively. Define*

$$\begin{aligned}\theta(B_i) &= B_{i+2}, \textit{for } 1 \le i \le 2\\ \varphi(b_i) &= b_{i+4}, \textit{for } 1 \le i \le 4\\ \sigma(a_i) &= a_{i+4}, \textit{for } 1 \le i \le 4\end{aligned}$$

Now, observe that

i) *for any red-yellow bigon* B_i *in L and any red edge* $a_i \in B_j$ *we have* $\sigma(a_i) = a_{i+4} \in \theta(B_j) = B_{j+2}$, *and*

ii) *for any blue edge* b_i *adjacent to a red edge* a_j *we have* $\varphi(b_i) = b_{i+4}$ *adjacent to* $\sigma(a_j) = a_{i+4}$.

Hence we conclude that K and L are congruent. As expected, a graph G underlies them both and this is the connected graph with two vertices and two edges and no loops. L models an imbedding of G on the sphere, $\mathbb{S}_0$, *and K models an imbedding of G on the non-orientable surface* $\mathbb{N}_1$. ❑

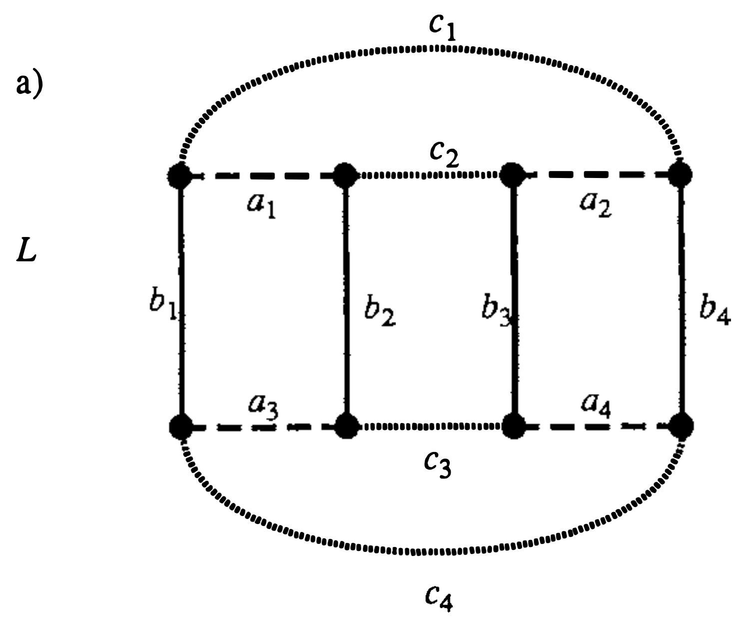

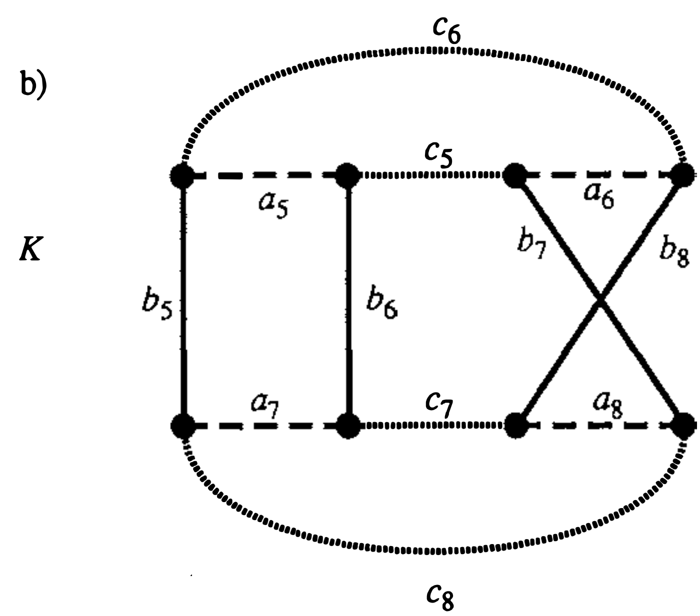

FIGURE IV.1

If K and L are congruent 3-graphs and $A = \{a_1, a_2, ..., a_n\}$ is a set of red edges in K, then for sake of conciseness we usually write A for $\sigma(A)$ and a_i for $\sigma(a_i)$ when no ambiguity results.

The following lemmas are immediate.

LEMMA IV.2. *If L is a 3-graph congruent to K then $c(L) = c(K)$, $|B(L)| = |B(K)|$, $|Y(L)| = |Y(K)|$, $|EL| = |EK|$, and $|VL| = |VK|$. Hence $\chi(L) = \chi(K) - |R(K)| + |R(L)|$.* ❑

LEMMA IV.3. *Let a be a red edge in a 3-graph K congruent to a 3-graph L. Let Y denote the red-blue (red-yellow) bigon in K that contains a. Let Y' denote the red-blue (red-yellow) bigon in L that contains a. Then $|Y| = |Y'|$.* ❑

2. THE GENUS AND CROSSCAP RANGES OF A 3-GRAPH

From Lemma II.11, any 3-graph K on $\mathbb{S}_g$ has genus

$$g(K) = 1 - \frac{\chi(K)}{2},$$

and any 3-graph K on $\mathbb{N}_k$ has crosscap number $k(K) = 2 - \chi(K)$.

The *genus range* of a 3-graph K is defined to be the set of numbers g such that there is a 3-graph in $\mathbb{S}_g$ congruent to K. The *minimum genus number* $g_{min}(K)$ of K is the minimum value in this range. The *maximum genus number* $g_{max}(K)$ of K is the maximum value in this range.

Similarly, the *crosscap range* is defined to be the set of numbers k such that there is a 3-graph in $\mathbb{N}_k$ congruent to K. The *minimum crosscap number* $k_{min}(K)$ of K is the minimum value in this range. The *maximum crosscap number* $k_{max}(K)$ of K is the maximum value in this range.

EXAMPLE IV.4. *Consider the* 3-*graphs* L *and* K *in Figure* IV.1. *No other orientable* 3-*graph is congruent to* L, *and therefore we conclude that the genus range for* L *is* $\{0\}$. *Likewise, no other non-orientable* 3-*graph is congruent to* L *other than* K, *and therefore we conclude that the crosscap range for* L *is* $\{1\}$. ❑

3. UNIPOLES AND POLES

Let v and w be a pair of adjacent vertices in a 3-graph K. Suppose that v and w are linked by a single edge a, which is red. Let b_1 and b_2 be the blue edges incident on v and w respectively. Let c_1 and c_2 be the yellow edges incident on v and w respectively. Let v_1, v_2, w_1, w_2 be the vertices other than v and w incident on c_1, b_1, c_2, b_2 respectively. Suppose that $\{c_1, b_1, c_2, b_2\}$ is included in a single blue-yellow bigon R. Then we say that a is a *red* 1-*unipole*. Furthermore a is *consistent* (Figure IV.2a) if $w_2 \in R_{v_1}[v_2, w_1]$; otherwise a is *inconsistent* (Figure IV.2b). Similar definitions can be made for *blue* and *yellow* 1-*unipoles*, but in this chapter it is only the red 1-unipoles that are of interest.

A red 1-*pole* is a red 1-dipole or a red 1-unipole and its *type* is *red*. Similar definitions can be made for *blue* and *yellow* 1-*poles*.

Hence an edge is a 1-pole unless it is a member of a digon. If L is congruent to K, then a 1-pole a in K is a 1-pole in L, by Lemma IV.3.

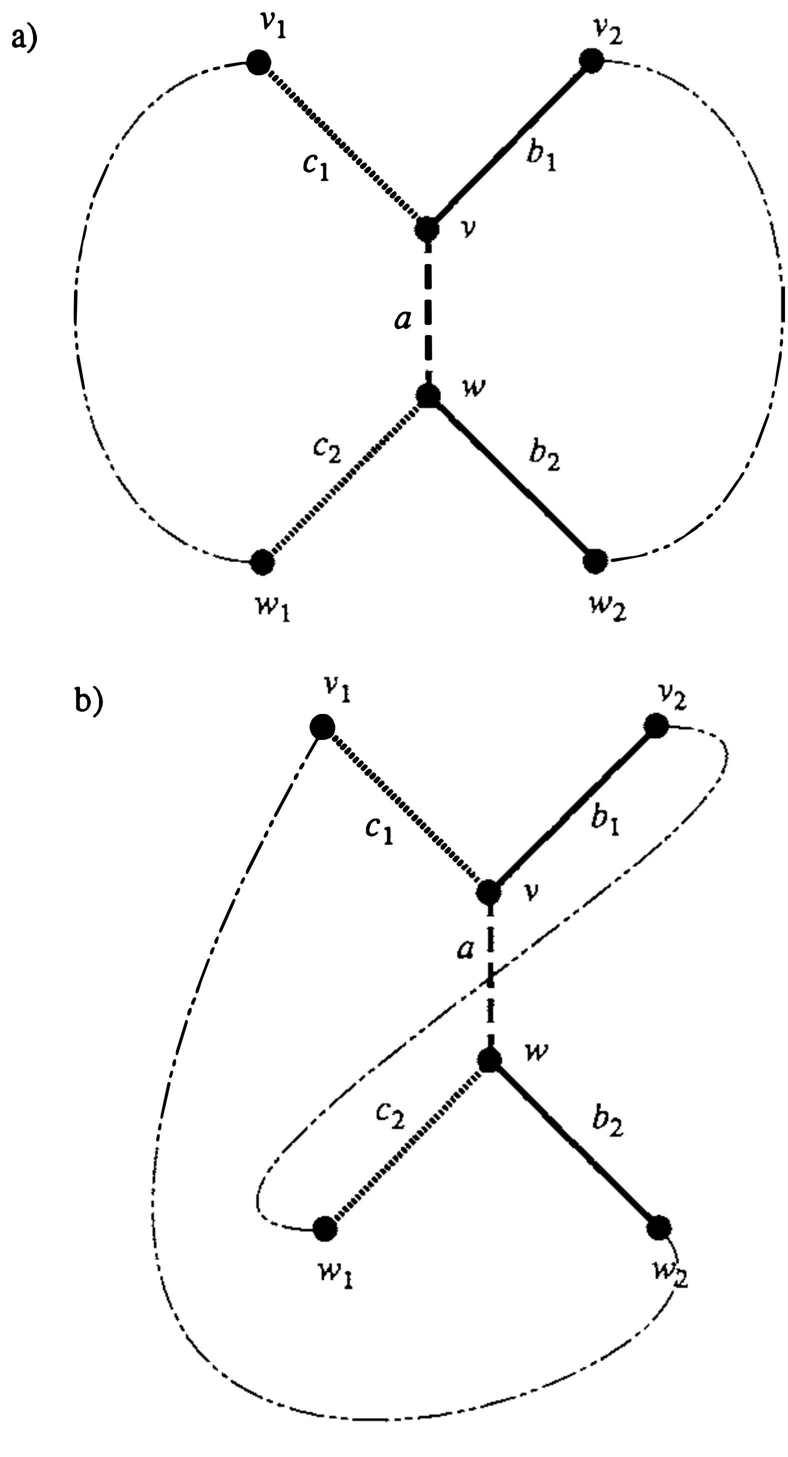

FIGURE IV.2

The *cancellation* of the 1-unipole a of Figure IV.2 is the operation of deletion of the vertices v and w followed by the insertion of edges c and b linking v_1 to w_1 and v_2 to w_2 respectively (Figure IV.3). We denote the resulting graph by $K - [a]$.

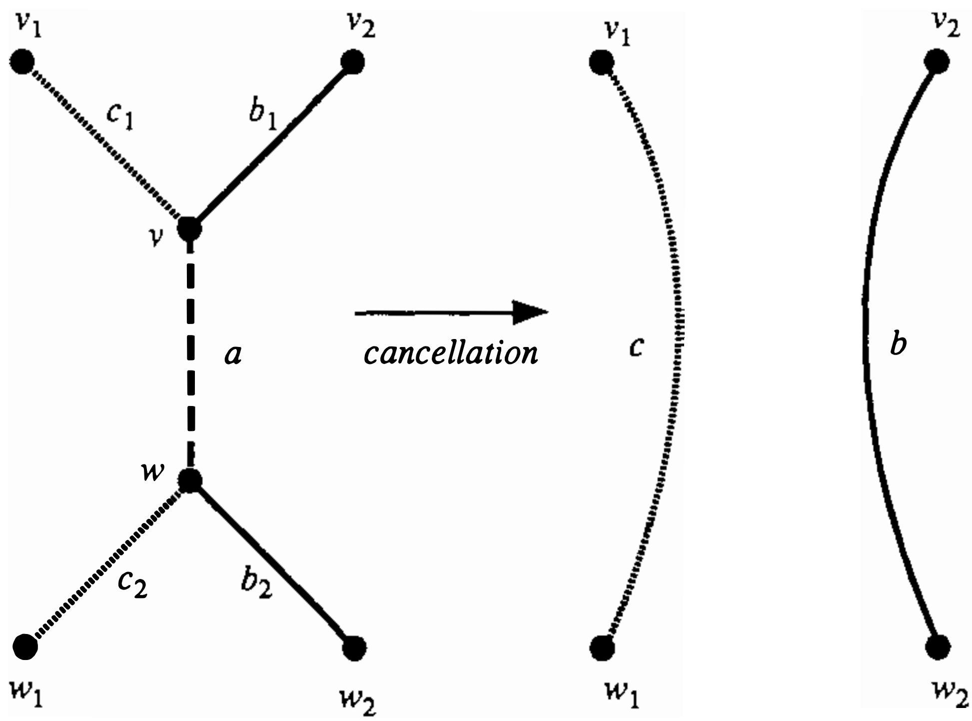

FIGURE IV.3

Now assume that K is bipartite. Therefore, one may colour the vertices of K black or white so that adjacent vertices receive distinct colours. Evidently v and w receive distinct colours, as do v_1 and w_1, and v_2 and w_2. Hence we conclude that $K - [a]$ is bipartite. This conclusion is best stated in the following way.

LEMMA IV.5. *If a is a 1-pole in an orientable 3-graph K, then $K - [a]$ is orientable.* ❑

The following lemma is immediate.

LEMMA IV.6. *If K and L are congruent 3-graphs and a is a 1-pole in K then $K - [a]$ and $L - [a]$ are congruent 3-graphs.* ❑

From Theorem II.2, we observed that cancellation of a 1-dipole does not alter the Euler characteristic of a 3-graph. However, Lemma IV.7 and IV.8 below show that the Euler characteristic is altered by the cancellation of a 1-unipole.

LEMMA IV.7. *If a is a consistent 1-unipole in a 3-graph K, then $\chi(K - [a]) = \chi(K) + 2$.*

Proof. The cancellation of a causes the number of vertices to drop by two. However, from Figure IV.2a, we see that the number of bigons increases by one. Hence

$$\begin{aligned} \chi(K - [a]) &= r(K) + 1 - \frac{|VK| - 2}{2} \\ &= \chi(K) + 2. \;\square \end{aligned}$$

LEMMA IV.8. *If a is an inconsistent 1-unipole in a 3-graph K, then $\chi(K - [a]) = \chi(K) + 1$. Furthermore $K - [a]$ is connected if K is connected.*

Proof. The cancellation of a causes the number of vertices to drop by two. However, from Figure IV.2b, we see that the number of bigons remains the same. Hence

$$\chi(K - [a]) = r(K) - \frac{|VK| - 2}{2}$$
$$= \chi(K) + 1.$$

Since there exists a path in $K - [a]$ joining w_2 to v_1 whose edges are included in $R - \{c_1, b_1, c_2, b_2\}$, we conclude that $K - [a]$ is connected if K is connected. ❑

$K - [a_5]$

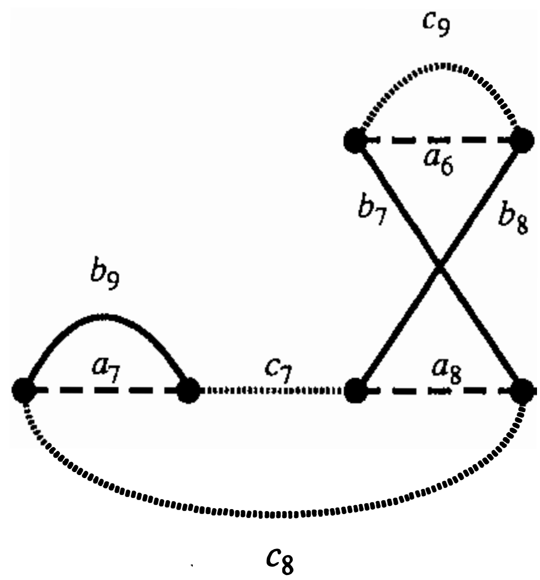

FIGURE IV.4

EXAMPLE IV.9. *Consider the* 3-*graph* K *in Figure* IV.1. *Since it has just one blue-yellow bigon, then all four red edges in* K *are* 1-*unipoles. It is easy to check that all four are in fact inconsistent. Figure* IV.4 *gives the planar* 3-*graph* $K - [a_5]$ *obtained by cancelling* a_5 *from* K. *We note in the resulting graph that* a_7 *is not a* 1-*pole. This is because the red-blue bigon in* K *that contains* a_5 *is a bisquare. Likewise,* a_6 *is not a* 1-*pole in* $K - [a_5]$. *However* a_8 *is a* *consistent* 1*-unipole in* $K - [a_5]$. *Cancellation of this edge results in a disconnected* 3*-graph.* ❑

We now give two situations where we can guarantee the consistency of a 1-unipole. Let a be a 1-pole in a 3-graph K. If $c(K - [a]) = c(K) + 1$ then we say that a is a *cut edge*.

LEMMA IV.10. *If a is a cut edge in a* 3*-graph* K, *then a is consistent. Hence* $\chi(K - [a]) = \chi(K) + 2$.

Proof. By Lemma IV.8 and the fact that cancelling a 1-dipole does not alter the number of components in a 3-graph, we conclude that a is a consistent 1-unipole. ❑

In Example IV.9, a_8 is a red cut edge in $K - [a_5]$.

LEMMA IV.11. *If a is a* 1*-unipole in an orientable* 3*-graph* K, *then a is consistent. Hence* $\chi(K - [a]) = \chi(K) + 2$.

Proof. Label the edges and vertices adjacent to a as in Figure IV.2b. If a were inconsistent, then $R_{v_1}[v, w] \cup \{a\}$ would be a circuit of odd length, contradicting the fact that K is bipartite. ❑

COROLLARY IV.12. *All* 1*-poles in a connected planar* 3*-graph are* 1*-dipoles or cut edges.*

Proof. Suppose we have a 1-unipole which is not a red cut edge in a connected planar 3-graph. Then by Lemma IV.11, cancellation of this unipole will result in a connected 3-graph with Euler characteristic 4, a contradiction to the fact that the Euler characteristic of any connected 3-graph is no more than 2. ❑

The only 1-pole in the 3-graph $K - [a_5]$ of Example IV.9 is the red edge a_8. It is a red cut edge as expected from Corollary IV.12.

4. REATTACHMENTS AND TWISTS

Let B be a red-yellow bigon in a 3-graph K. Let a be a red 1-pole in B. Furthermore, suppose $c \in B$ to be a yellow edge that is not adjacent to a. The following uses the notation of Figure IV.5. Let $K' = K - [a]$ and let b denote the blue edge in K' joining v_1 to v_2. Let c' be the yellow edge in K' joining v_3 to v_4. (Figure IV.5b.) Consider the following two cases.

Case (a) Let L be the graph with vertex set $VK' \cup \{v_7, v_8\}$, where $v_7, v_8 \notin VK'$, and edge set $(EK' - \{b, c\}) \cup \{b_1, b_2, a', c_1, c_2\}$ such that $\psi b_i = \{v_i, v_{i+6}\}$, $\psi c_i = \{v_{i+4}, v_{i+6}\}$ and $\psi a' = \{v_7, v_8\}$. Colour each b_i blue, each c_i yellow and a' red. (Figure IV.5c.) We say that L is obtained from K by a *reattachment* of a to c. The colouring of the vertices in the figure demonstrates that L is orientable if and only if K is orientable. Furthermore L is congruent to K and connected if K is connected. Reattaching a' to c' yields the original 3-graph K, and hence the operation of reattachment is reversible.

Case (b) Let L be the graph with vertex set $VK' \cup \{v_7, v_8\}$ where $v_7, v_8 \notin VK'$, and edge set $(E(K') - \{b, c'\}) \cup \{b_1, b_2, a', c_1, c_2\}$ such that $\psi b_i = \{v_i, v_{i+6}\}$, $\psi c_i = \{v_{i+2}, v_{i+6}\}$ and $\psi a' = \{v_7, v_8\}$. Colour each b_i blue, each c_i yellow and a' red. (Figure IV.5d.) We say that L is the 3-graph obtained from K by a *twist* of a. Clearly L is congruent to K and connected if K is connected. Twisting a' yields the original 3-graph K, and hence the twist operation is reversible. However, orientation may not be preserved during a twist, as shown in Lemmas IV.14, and IV.16.

a)

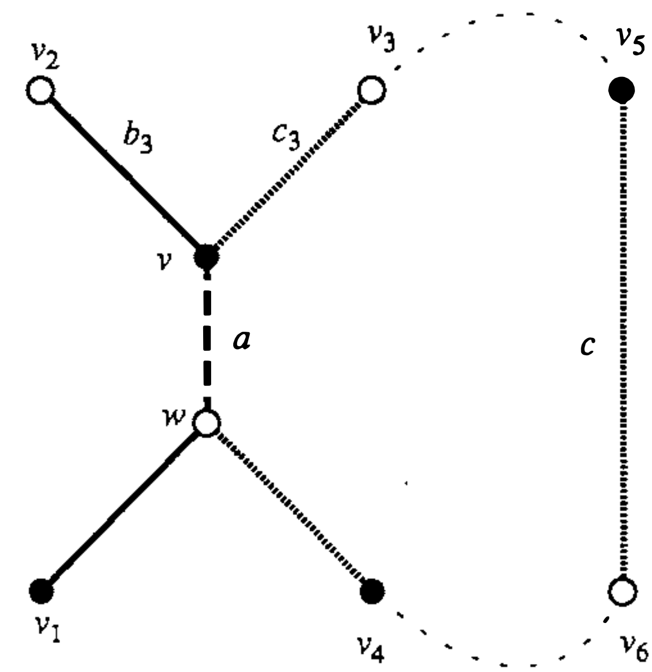

b)

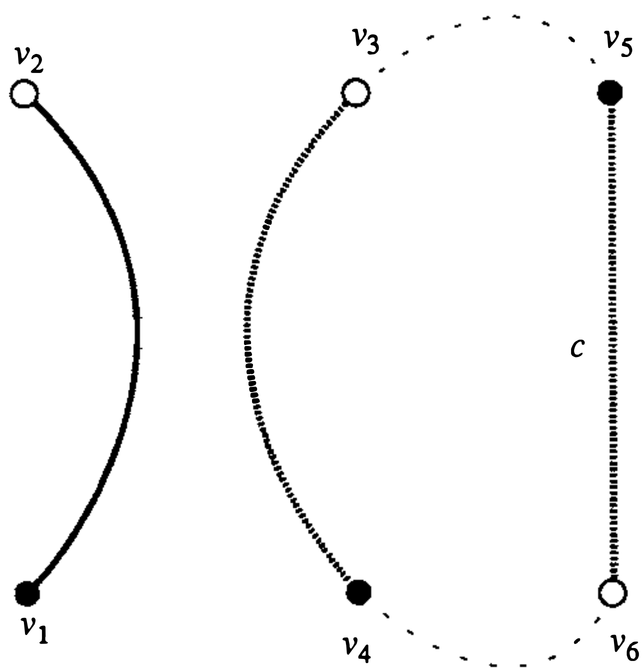

c)

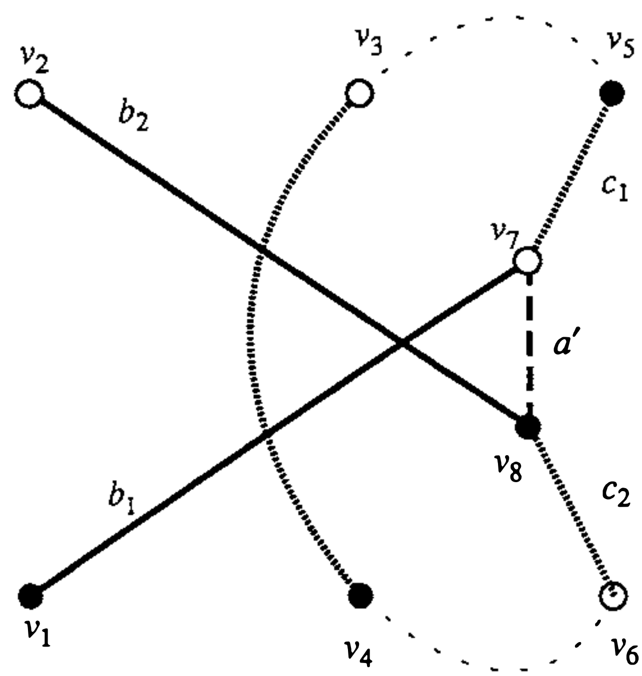

d)

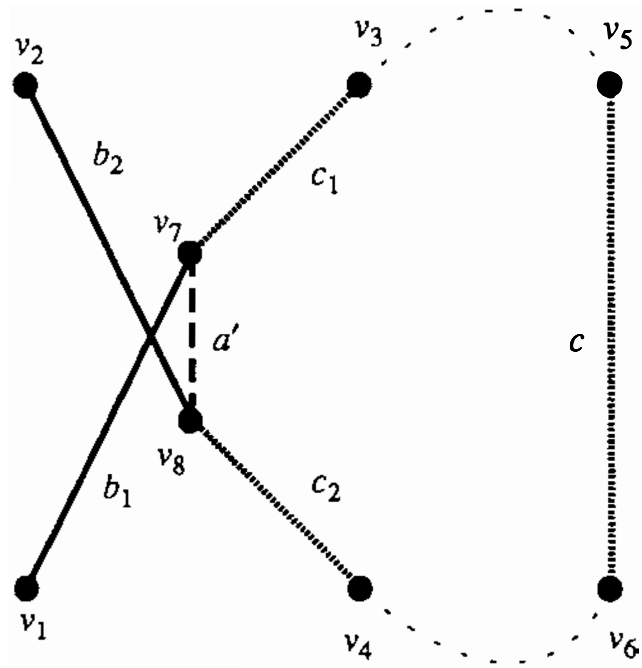

FIGURE IV.5

EXAMPLE IV.13. *Consider the orientable* 3*-graphs* K *and* L *in Figure* IV.6. *Clearly* L *is obtained from* K *by reattaching* a *to* c. ❑

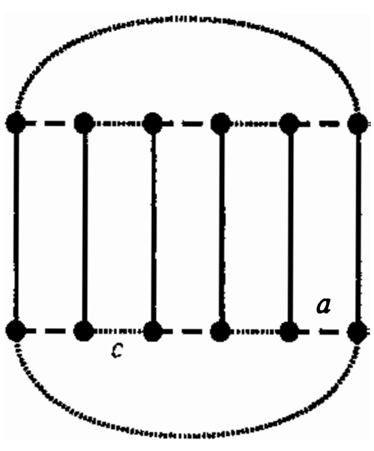

K

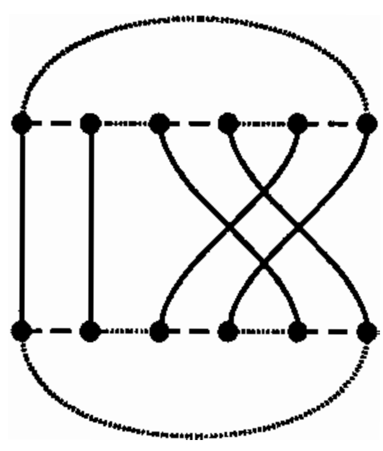

L

FIGURE IV.6

LEMMA IV.14. *Let a be a red* 1*-dipole in a (connected)* 3*-graph K. Then the (connected)* 3*-graph L obtained from K by a twist of a is non-orientable.*

Proof. The following uses the notation of Figure IV.5. Let R denote the blue-yellow bigon in K that passes through v. Since a is a 1-dipole, $VR \cap \{v_1, v_4\} = \varnothing$. Hence in L, $R_v[v_2, v_3] \cup \{c_1, a', b_2\}$ is a circuit of odd length. Hence we conclude that L is non-orientable. ❑

EXAMPLE IV.15. *Consider the* 3*-graphs L and K in Figure* IV.1. *By a twist of the* 1*-pole* a_2 *in the orientable* 3*-graph L we obtain the non-orientable* 3*-graph K.* ❑

LEMMA IV.16. *Let a be a red* 1*-unipole in an orientable* 3*-graph K. Then the* 3*-graph L obtained from K by a twist of a is orientable if and only if a is a red cut edge.*

Proof. We may assume K, and hence L, connected. The following uses the notation of Figure IV.5.

Firstly suppose a is not a red cut edge. Hence $K - [a]$ is connected and bipartite, since K is. Let R_1 and R_2 be the blue-yellow bigons in $K - [a]$ that contain b and c' respectively. Consider the following cases.

Case i) if $R_1 = R_2$ then clearly R_1 includes a path P that joins v_2 to v_3. Moreover $b, c' \notin P$, for otherwise K would not be bipartite.

Case ii) if R_1 and R_2 are distinct, then since $K - [a]$ is connected, there exists a path P from v_2 to v_3 in $K - [a]$. Furthermore, $P + R_1$ and $P + R_2$ both contain paths that join v_2 to v_3 in $K - [a]$. Therefore we may choose P so that $b, c' \notin P$.

In both cases P is a path in K. Furthermore, $P \cup \{b_3, c_3\}$ is a circuit in K, which must be even since K is bipartite. Hence $P \cup \{c_1, a', b_2\}$ is a circuit of odd length in L, and we conclude that L is non-orientable.

Now suppose a is a red cut edge. Let L_1 and L_2 be the two components of $K - [a]$ where $c' \in EL_1$ and $b \in EL_2$. Since K is orientable, L_1 and L_2 are orientable. We may colour the vertices of L_1 and L_2 black and white so that no two vertices of the same colour are adjacent. Furthermore, we may assume without loss of generality that v_1 and v_3 are coloured white. Therefore v_2 and v_4 are coloured black. It is immediate that L is bipartite since colouring v_7 black and v_8 white gives a colouring of VL in two colours so that no two vertices of the same colour are adjacent. Hence L is orientable. ❑

LEMMA IV.17. *Let a be a red* 1-*dipole in a* 3-*graph K. Let L denote the non-orientable* 3-*graph obtained from K by a twist of a. Then* $\chi(L) = \chi(K) - 1$ *and hence* $k(L) = k(K) + 1$.

Proof. Clearly $|R(L)| = |R(K)| - 1$. Since L and K are congruent then by Lemma IV.2 we have $\chi(L) = \chi(K) - |R(K)| + |R(L)| = \chi(K) - 1$. ❑

5. THE BETTI NUMBER OF A 3-GRAPH

The *Betti number* of a connected 3-graph K, denoted $b(K)$, is defined by the equation $b(K) = 1 - |B(K)| - |Y(K)| + |EK| - |VK| = 1 - \chi(K) + |R(K)|$. For example, the 3-graph L in Figure IV.1 has Betti number $b(L) = 1 - 2 + 2 = 1$. This is also the maximum crosscap number of a non-orientable 3-graph congruent to L. The following theorem establishes that this observation occurs generally.

THEOREM IV.18. *Let K be a connected 3-graph. Then* $k_{max}(K) = b(K)$.

Proof. Let L be a 3-graph congruent to K with the minimum number of blue-yellow bigons. Hence $k_{max}(K) = k(L) = 2 - \chi(L)$. By Lemma IV.2, we have

$$\begin{aligned} k_{max}(K) &= 2 - \chi(L) \\ &= 2 - \chi(K) + |R(K)| - |R(L)| \\ &= 1 + b(K) - |R(L)|. \end{aligned}$$

We claim $|R(L)| = 1$, and hence $k_{max}(K) = 1 + b(K) - 1 = b(K)$. To see this claim, we suppose L to have at least two blue-yellow bigons. By the connectedness of L, there exists a red 1-dipole a in L. By Lemma IV.14, the 3-graph L' obtained from K by twisting a is non-orientable. Hence by Lemma IV.17, L' is a 3-graph, congruent to K, such that $k(L') = k(L) + 1 = k_{max}(K) + 1$, a contradiction. ❑

THEOREM IV.19. *Let K be a connected 3-graph. Then the crosscap range is an unbroken interval of integers.*

Proof. Let L_0 be a non-orientable 3-graph congruent to K with crosscap number $k(L_0) = k_{min}(K)$. If L_0 contains just one blue-yellow bigon then evidently $k_{max}(K) = k_{min}(K)$, and hence we are done. Now assume that L_0 contains at least two blue-yellow bigons. Evidently L_0 contains a red 1-dipole. Let L_1 be the 3-graph obtained from L_0 by twisting a red 1-dipole. By Lemma IV.14, L_1 is a connected non-orientable 3-graph congruent to K. Moreover, by Lemma IV.17, $k(L_1) = k(L_0) + 1$. Proceeding inductively for $k_{max}(K) - k_{min}(K)$ steps, we obtain a sequence $L_0, L_1, ..., L_n$ of non-orientable 3-graphs such that $k(L_i) = k(L_{i+1}) - 1$ whenever $0 \leq i < n$. Furthermore L_n has just one blue-yellow bigon and hence $k(K_n) = k_{max}(K)$, as required. ❑

6. PERMITTED POLE SETS

LEMMA IV.20. *Suppose $\{a, a', e_1, e_2\}$ is a bigon in a 3-graph K, and suppose that a and a' are 1-poles of the same type. Then a is not a 1-pole in $K - [a']$.*

Proof. Clearly e_1 and e_2 coalesce to form an edge e' in $K - [a']$, which implies that $\{a, e'\}$ is a digon in $K - [a']$. Hence a is not a 1-pole in $K - [a']$. ❑

LEMMA IV.21. *Suppose a and a' are 1-poles of the same type in a 3-graph K, and that a is a 1-pole in $K - [a']$. Then a' is a 1-pole in $K - [a]$.*

Proof. Suppose that a' is not a 1-pole in $K - [a]$. Then there exists a digon $\{a', e\}$ in $K - [a]$. Since a' is a 1-pole in K, then $e \notin EK$, and hence there exist edges e_1 and e_2 in K that coalesce to form e in $K - [a]$. Futhermore both e_1 and e_2 are adjacent to a and a'. Therefore $\{a, a', e_1, e_2\}$ is a bigon in K. By Lemma IV.20, a is not a 1-pole in $K - [a']$, a contradiction. ❑

Lemma IV.21 motivates the following definition. A set $A = \{a_1, a_2, ..., a_n\}$ of red edges in a 3-graph K is a *permitted red pole set* if a_1 is a 1-pole in K, and each a_i is a 1-pole in $K - [a_1] - [a_2] - ... - [a_{i-1}]$ whenever $2 \leq i \leq n$. We usually write $K - [A]$ or $K - [a_1, a_2, ..., a_n]$ for $K - [a_1] - [a_2] - ... - [a_n]$.

LEMMA IV.22. *Suppose a and a' are 1-poles of the same type in a 3-graph K, and that a is a 1-pole in $K - [a']$. Then $K - [a] - [a']$ $= K - [a'] - [a]$.*

Proof. By Lemma IV.21, $K - [a] - [a']$ is defined. Consider the following cases.

i) no edge in K is adjacent to both a and a'. The following uses the notation of Figure IV.7a. Clearly $c_1, c_2, c_3, c_4, b_1, b_2, b_3$ and b_4 are distinct edges. Let L be the 3-graph obtained from K by deleting the edges $a, a', c_1, c_2, c_3, c_4, b_1, b_2, b_3, b_4$ and inserting yellow edges c and c' and blue edges b and b' such that $\psi c = \{v_1, v_2\}$, $\psi c' = \{v_3, v_4\}$, $\psi b = \{v_5, v_6\}$ and $\psi b' = \{v_7, v_8\}$. (See Figure IV.7b.)

ii) a and a' are both adjacent to a single edge b'. Without loss of generality assume that b' is a blue edge. The following uses the notation of Figure IV.7c. Clearly $c_1, c_2, c_3, c_4, b_1, b'$ and b_2 are distinct edges. Let L be the 3-graph obtained from K by deleting the edges $a, a', c_1, c_2, c_3, c_4, b_1, b', b_2$ and inserting yellow edges c and d and a blue edge b such that $\psi c = \{v_1, v_2\}$, $\psi d = \{v_3, v_4\}$ and $\psi b = \{v_5, v_6\}$. (See Figure IV.7d.)

iii) a and a' are both adjacent to a yellow edge c' and a blue edge b'. The two possible situations are shown in Figures IV.7e and f. Clearly c_1, c', c_2, b_1, b' and b_2 are distinct edges. Let L be the 3-graph obtained from K by deleting the edges $a, a', c_1, c', c_2, b_1, b', b_2$ and inserting a yellow edge c and a blue edge b such that $\psi c = \{v_1, v_2\}$ and $\psi b = \{v_3, v_4\}$. (See Figure IV.7g.)

No other cases are possible, for otherwise either a or a' would not be a 1-pole. In all cases, evidently $L = K - [a] - [a'] = K - [a'] - [a]$, as required. ❑

a)

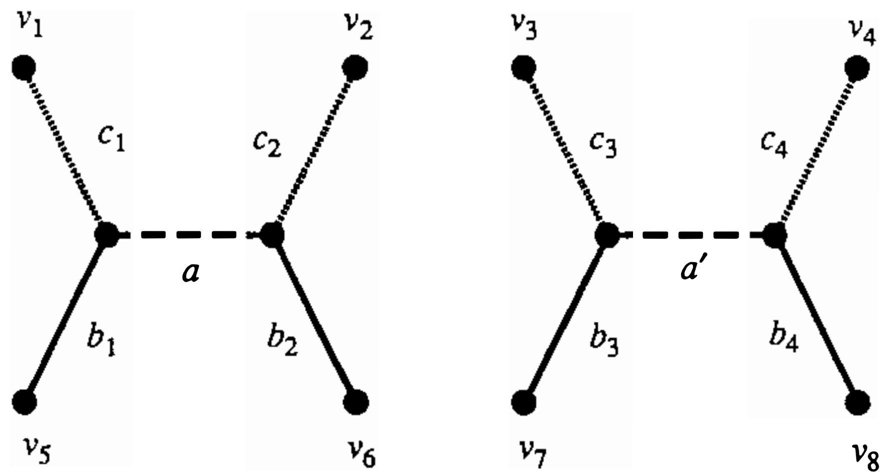
v_1
v_2
v_3
v_4
c_1
c_2
c_3
c_4
a
a'
b_1
b_2
b_3
b_4
v_5
v_6
v_7
v_8

b)

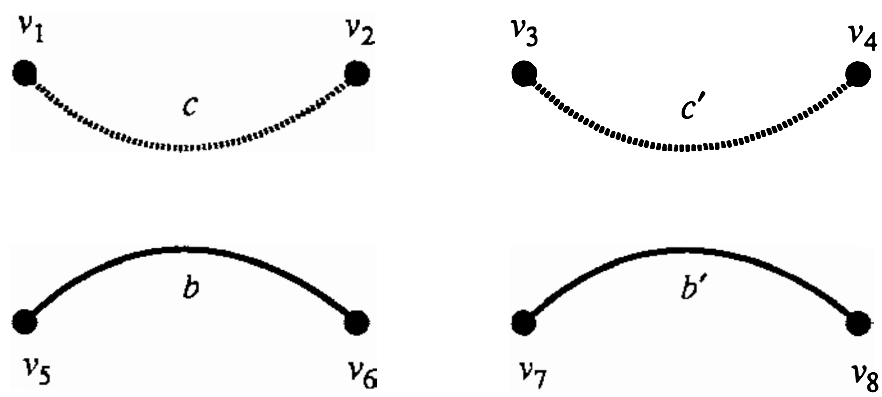
v_1
v_2
v_3
v_4
c
c'
b
b'
v_5
v_6
v_7
v_8

c)

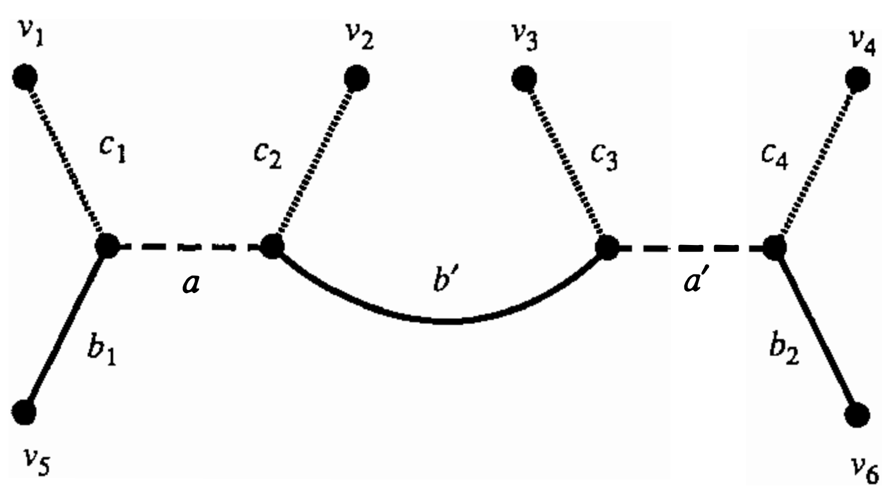
v_1
v_2
v_3
v_4
c_1
c_2
c_3
c_4
a
b'
a'
b_1
b_2
v_5
v_6

d)

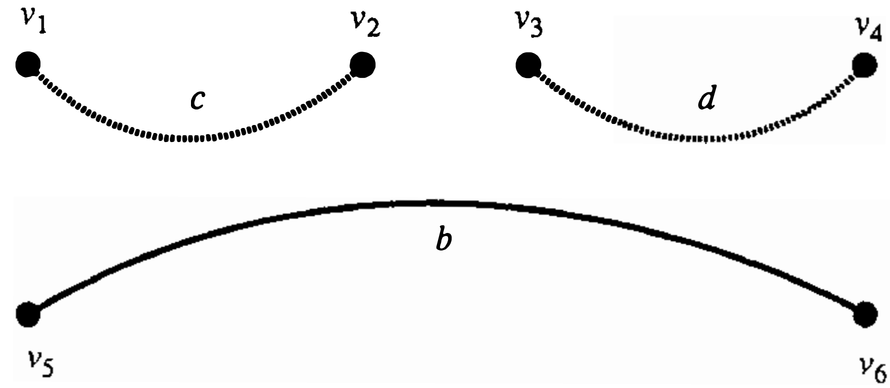
v_1
v_2
v_3
v_4
c
d
b
v_5
v_6

e)

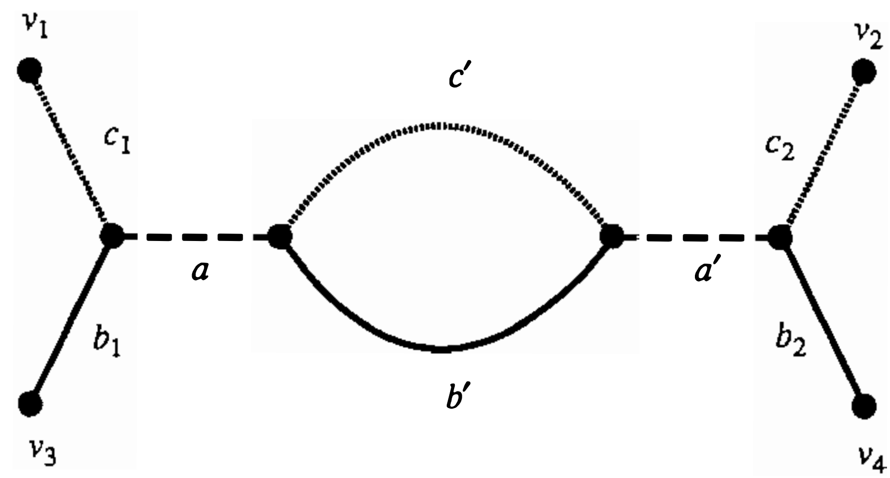
v_1
v_2
c'
c_1
c_2
a
a'
b_1
b_2
b'
v_3
v_4

f)

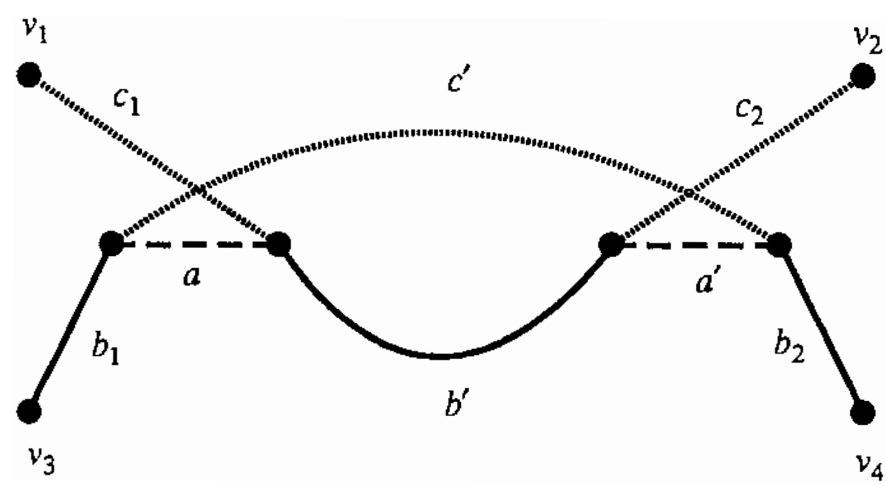
v_1
v_2
c_1
c'
c_2
a
a'
b_1
b_2
b'
v_3
v_4

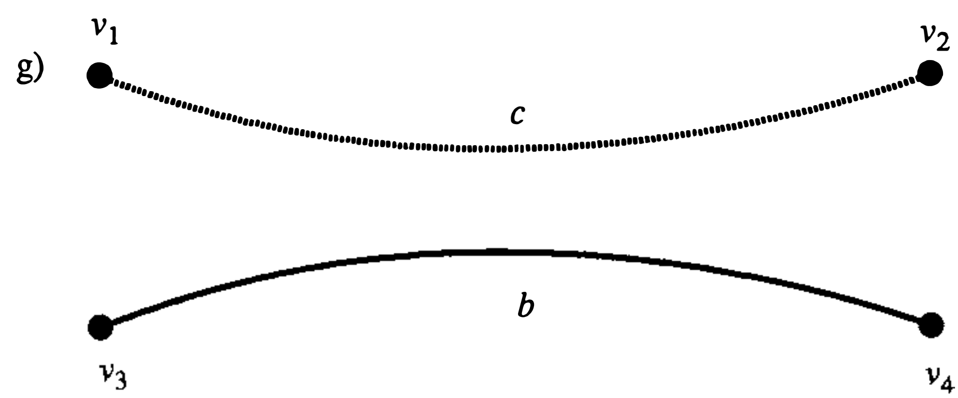

FIGURE IV.7

COROLLARY IV.23. *Let A be a permitted red pole set in a 3-graph K, and let B be obtained from A by permuting the elements of A. Then B is a permitted red pole set. Furthermore $K - [A] = K - [B]$.*

Proof. This follows directly from Lemmas IV.21 and IV.22, and the fact that a permutation of a set can be obtained by a sequence of transpositions. ❑

From Lemma IV.21, we conclude that the order in which we cancel the a_i's does not matter, and hence we are justified in using the term "set" rather than "sequence".

A permitted red pole set A in a 3-graph K is *maximal* if $K - [A]$ contains no red 1-pole. If L is congruent to K, then by Lemma IV.3, $L - [A]$ contains no red 1-pole when A is maximal in K. Hence A is maximal in L.

EXAMPLE IV.24. *Consider the 3-graphs L and K of Example* IV.1. *The two maximal permitted red pole sets in L are* $\{a_1, a_4\}$ *and* $\{a_2, a_3\}$. *Clearly* $\sigma\{a_1, a_4\}$ *and* $\sigma\{a_2, a_3\}$ *are the two maximal permitted red pole sets in K.* ❑

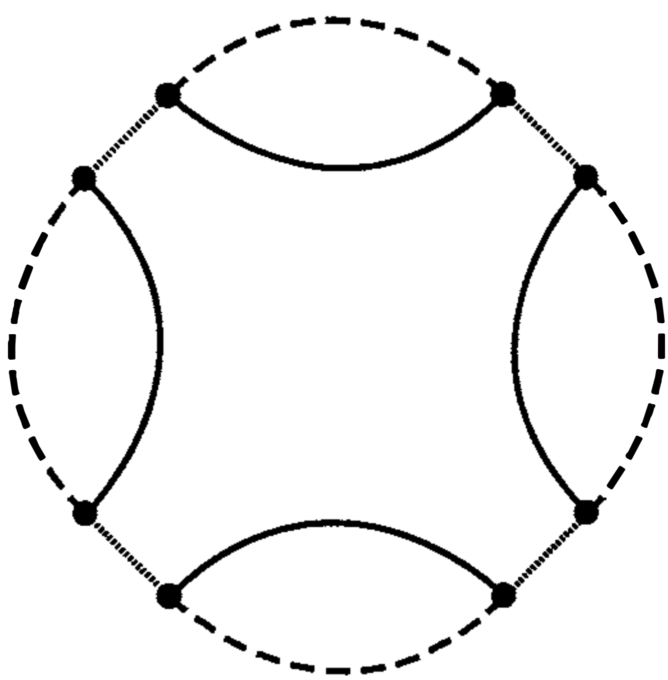

FIGURE IV.8

7. RINGS

A *yellow n-ring* is a connected 3-graph with just one blue-yellow bigon and n red-blue bigons. Furthermore each red-blue bigon is a digon. Similarly, a *blue n-ring* is a connected 3-graph with just one blue-yellow bigon and n red-yellow bigons. Furthermore each red-yellow bigon is a digon. A *ring* is a yellow or blue n-ring. Since interchanging the red and blue edges in a yellow ring results in an isomorphic 3-graph, then we conclude that there is one red-yellow bigon in a yellow ring. Similarly, there is one red-blue bigon in a blue ring.

Figure IV.8 gives an example of a yellow 4-ring.

It is clear that two yellow n-rings are isomorphic 3-graphs as are two blue n-rings. Hence we talk about "the" yellow n-ring and "the" blue n-ring for some positive integer n. For example, the 1-ring is the trivial 3-graph. However, the next lemma states a stronger result. That is, any 3-graph congruent to a ring is in fact that ring.

LEMMA IV.25. *Let K and L be two congruent* 3-*graphs such that each component in L is a ring. Then K and L are isomorphic* 3-*graphs.*

Proof. We may assume both K and L connected, as the general result follows by treating each component separately. Therefore L is a ring.

Firstly assume that L is a yellow n-ring, for some integer n. Hence L has just one red-yellow bigon. By congruence, we have that K has just one red-yellow bigon. Furthermore we have a bijection between the red-blue bigons of L and the red-blue bigons of K. Since each red-blue bigon in L is a digon, then by Lemma IV.3 each red-blue bigon in K is a digon. This immediately implies that K is a yellow n-ring, isomorphic to L, as required.

One can obtain a similar result if L is a blue n-ring. ❑

LEMMA IV.26. *If A denotes a maximal permitted red pole set in a* 3*-graph K, then each component of $K - [A]$ is a ring.*

Proof. Let K' denote a component of $K - [A]$. If K' contains at least two blue-yellow bigons, then there is a red 1-dipole in K', a contradiction.

Let a denote a red edge in K'. The following uses the notation of Figure IV.2. Since a is not a 1-unipole or a 1-dipole, either $b_1 = b_2$ or $c_1 = c_2$. Firstly, assume that $b_1 = b_2$. If $c_1 = c_2$ then K' is the 1-ring and we are done. Hence assume that $c_1 \neq c_2$. Let a' denote the red edge other than a adjacent to c_1. Clearly $a \neq a'$ and hence two distinct yellow edges are adjacent to a'. Therefore a', like a, is a member of a red-blue digon. Proceeding inductively we conclude that all red edges in K' belong to red-blue digons. Hence K' is a yellow ring.

Similarly, if $b_1 \neq b_2$ and $c_1 = c_2$ then K' is a blue ring. ❑

EXAMPLE IV.27. *Consider the* 3*-graph K in Example* IV.1. *In Example* IV.24, *we saw that $A = \{a_1, a_3\}$ is a maximal permitted red pole set. Figure* IV.9 *illustrates the 3-graphs $K - [a_1]$ and $K - [a_1, a_3] = K - [A]$. We note that the two components of $K - [A]$ are yellow (or red)* 1*-rings, as expected from Lemma* IV.26. ❑

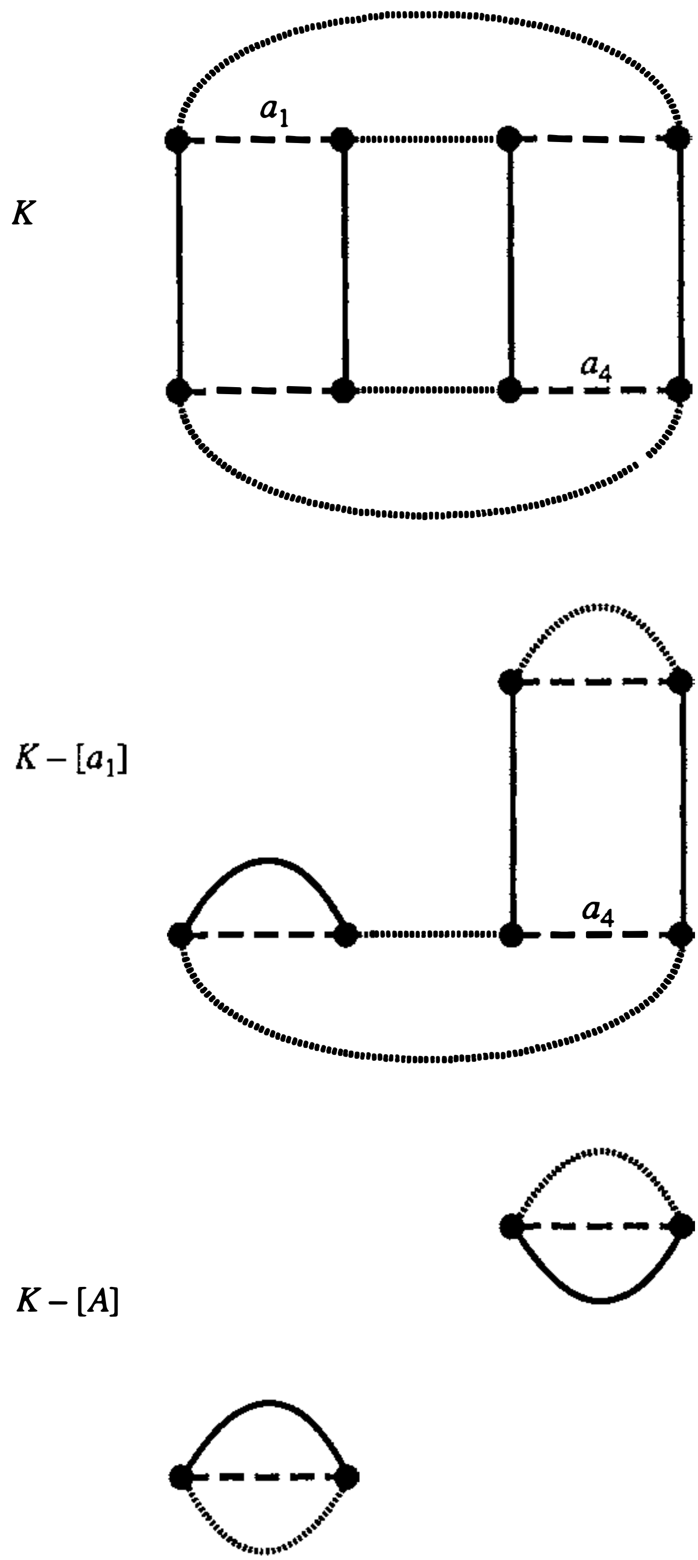

FIGURE IV.9

8. ζ-MOVES

A *ζ-move* (on a) is a reattachment of a red 1-pole a or a twist of a red cut edge a. Two 3-graphs, K and L, are *ζ-equivalent* if K can be obtained from L by a finite sequence of ζ-moves. Furthermore, K and L are *ζ-adjacent* if K can be obtained from L by one ζ-move.

It is immediate that two ζ-equivalent 3-graphs are congruent, the inverse of a ζ-move is a ζ-move, and ζ-moves preserve connectedness. We now show that ζ-moves preserve orientability. This is why we restrict our twists to red cut edges.

LEMMA IV.28. *Let K and L be two ζ-equivalent 3-graphs. Then K is orientable if and only if L is orientable.*

Proof. Clearly we may assume that K and L are ζ-adjacent. If the ζ-move in question is a reattachment, then our lemma follows from the observation, following the definition of reattachment, that K is bipartite if and only if L is bipartite.

Now suppose that the ζ-move in question is a twist of a red cut edge a. By Lemma IV.16, K is orientable if and only if L is orientable, as required. ❑

LEMMA IV.29. *Let K and L be two orientable ζ-adjacent 3-graphs. Then $|\chi(K) - \chi(L)| = 0$ or 2. Hence $|g(K) - g(L)| = 0$ or 1.*

Proof. Suppose a is a red 1-pole in K and L is obtained from K by a reattachment of a. Firstly, assume that a is a red 1-dipole.

Then by Theorem II.3, $\chi(K - [a]) = \chi(K)$. Since $K - [a] = L - [a]$, then $\chi(L) = \chi(K)$ if a is a 1-dipole in L or $\chi(L) = \chi(K) - 2$ if a is a 1-unipole in L, by Lemma IV.11. Now, assume that a is a red 1-unipole. Then by Lemma IV.11, $\chi(K - [a]) = \chi(K) + 2$. Since $K - [a] = L - [a]$, then $\chi(L) = \chi(K) + 2$ if a is a 1-dipole in L or $\chi(L) = \chi(K)$ if a is a 1-unipole in L, by Lemma IV.11. Hence $|\chi(K) - \chi(L)| = 0$ or 2.

Now suppose that a is a red cut edge in K and L is obtained from K by a twist of a. Clearly a is a cut edge in L, since $c(L - [a]) = c(K - [a]) = c(K) + 1 = c(L) + 1$. Hence by Lemma IV.10, $\chi(L) = \chi(L - [a]) - 2 = \chi(K - [a]) - 2 = \chi(K)$. ❑

9. THE EQUIVALENCE OF CONGRUENCE AND ζ-EQUIVALENCE

The following lemma follows from the definition of a reattachment of a red 1-pole or a red cut edge twist.

LEMMA IV.30. *Let L and K be two orientable congruent 3-graphs and let a be a red 1-pole in K. If $K - [a] = L - [a]$ then L can be obtained from K by a ζ-move on a. Hence K and L are ζ-adjacent.* ❑

LEMMA IV.31. *Let L and K be two orientable congruent 3-graphs and let a be a red 1-pole in K. If $K - [a]$ and $L - [a]$ are ζ-adjacent, then K and L are ζ-equivalent.*

Proof. The following uses the notation of Figure IV.3. Firstly, suppose $L - [a]$ is obtained from $K - [a]$ by a reattachment of a red 1-pole a' to a yellow edge c'. If $c' \neq c$, then $a', c' \in EK$, and we let K' denote the 3-graph obtained from K by reattaching a' to c'. If $c' = c$, we let K' denote the 3-graph obtained from K by reattaching a' to c_1. Therefore, K and K' are ζ-adjacent. In both cases we have $K' - [a] = L - [a]$, and hence by Lemma IV.30, K' and L are ζ-adjacent. Evidently, K and L are ζ-equivalent.

Now suppose that $L - [a]$ is obtained from $K - [a]$ by a twist of a red cut edge a'. Let K' denote the 3-graph obtained from K by twisting a'. Therefore, K and K' are ζ-adjacent. We have $K' - [a] = L - [a]$, and hence by Lemma IV.30, K' and L are ζ-adjacent. Evidently, K and L are ζ-equivalent, as required. ❑

We apply Lemma IV.31 inductively to obtain the following lemma.

LEMMA IV.32. *Let L and K be two orientable congruent 3-graphs and let a be a red 1-pole in K. If* $K - [a]$ *and* $L - [a]$ *are* ζ*-equivalent, then K and L are* ζ*-equivalent.* ❑

THEOREM IV.33. *Two orientable 3-graphs are congruent if and only if they are* ζ*-equivalent.*

Proof. The fact that two orientable ζ-equivalent 3-graphs are congruent follows immediately from our observations following the definition of ζ-equivalence.

Let J and L be two congruent orientable 3-graphs. We are required to show that J can be obtained from L by a finite sequence of red edge reattachments and red cut edge twists. Let $A = \{a_{n-1}, ..., a_2, a_1\}$ be a maximal permitted red pole set of L. Hence A is a maximal permitted red pole set in J.

Let $L_i = L - [a_{n-1}, ..., a_i]$ and $J_i = J - [a_{n-1}, ..., a_i]$ whenever $1 \leq i \leq n - 1$. By Lemma IV.6, L_i is congruent to J_i for each i. Hence $L_1 = L - [A]$ is congruent to $J_1 = J - [A]$. Let $J_n = J$ and $L_n = L$. We shall show by induction on i that J_n and L_n are ζ-equivalent. By Lemmas IV.26 and IV.25, $L_1 = J_1$, and it is immediate that L_1 and J_1 are ζ-equivalent.

Now assume that $L_{i-1} = L_i - [a_{i-1}]$ and $J_{i-1} = J_i - [a_{i-1}]$ are ζ-equivalent for some i such that $2 \leq i < n$. To complete the induction, we apply Lemma IV.32 and conclude that L_i and J_i are ζ-equivalent. ❑

EXAMPLE IV.34. *Consider the two congruent orientable 3-graphs K and L in Figures* IV.10*a and* IV.10*b respectively. By the previous theorem K and L are ζ-equivalent. We shall show that L can be obtained from K by two ζ-moves. Firstly, reattach a_2 to c_1. Figure* IV.10*c illustrates the resulting* 3-*graph. Then reattach a_1 to c_2. L is the resulting* 3-*graph.* ❑

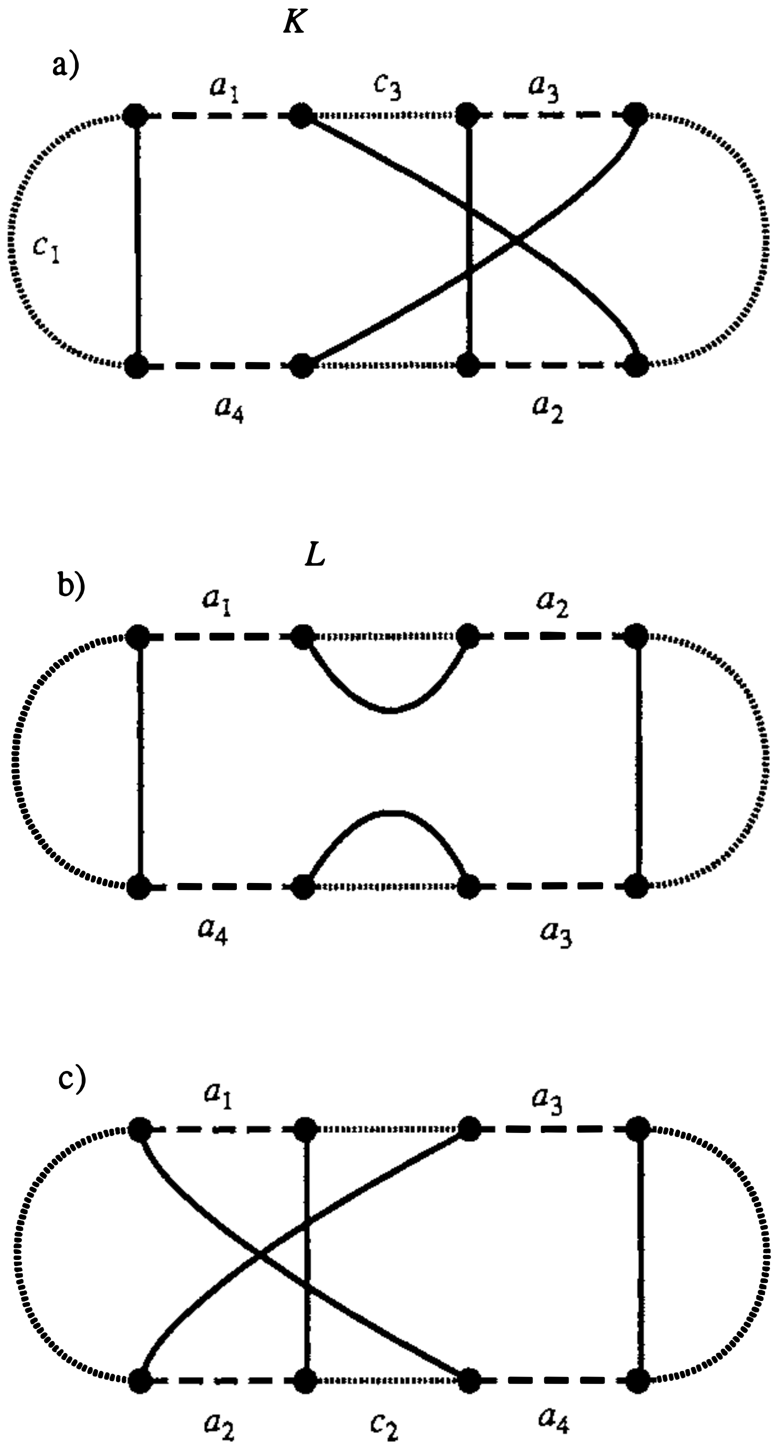

FIGURE IV.10

10. ORIENTABLE INTERPOLATION THEOREM

THEOREM IV.35. *Let K be a 3-graph. Then the genus range is an unbroken interval of integers.*

Proof. Let J and L be two orientable 3-graphs congruent to K. Since J and L are congruent then by Theorem IV.33 J may be obtained from L by a finite sequence of ζ-moves. Thus, there exists a sequence of ζ-adjacent 3-graphs congruent to K starting with one in the surface of genus $g_{min}(K)$ and ending with one in the surface of genus $g_{max}(K)$. By Lemma IV.29, the genera of the adjacent 3-graphs differ by at most one. The theorem follows. ❑

11. AN UPPER BOUND ON THE MINIMUM CROSSCAP NUMBER

Let G be a connected graph. The *genus range* of G is defined to be the set of numbers g such that G underlies a gem in $\mathbb{S}_g$. The *minimum genus number* $g_{min}(G)$ of G is the minimum value in this range. The *maximum genus number* $g_{max}(G)$ of G is the maximum value in this range.

Similarly, the *crosscap range* is defined to be the set of numbers k such that G underlies a gem in $\mathbb{N}_k$. The *minimum crosscap number* $k_{min}(G)$ of G is the minimum value in this range. The *maximum crosscap number* $k_{max}(G)$ of G is the maximum value in this range.

For a connected graph G, it is well known that $k_{min}(G) \leq 2g_{min}(G) + 1$. (See, for example, [6] pp. 136.) Theorem IV.36 below generalises this result to 3-graphs.

THEOREM IV.36. *Let K be a connected 3-graph. Then* $k_{min}(K) \leq 2g_{min}(K) + 1$.

Proof. Let L be a 3-graph congruent to K with genus $g(L) = g_{min}(K)$. If every red 1-pole in L is a cut edge, then clearly $k_{min}(K) = k(L) = g(L) = 0$.

Now suppose there exists a red 1-pole a in L that is not a cut edge. Let J be the 3-graph obtained from L by a twist of a. By Lemma IV.14, J is non-orientable. Consider the following two cases.

a) a is a red 1-dipole. Then by Lemma IV.17, $\chi(J) = \chi(L) - 1$. Therefore $k_{min}(K) \leq k(J) = 2 - \chi(J) = 2 - \chi(L) + 1 = 2 - 2 + 2g(L) + 1 = 2g_{min}(K) + 1$, as required.

b) a is a red 1-unipole. Then $r(J) = r(L)$ and therefore $\chi(J) = \chi(L)$. Therefore $k_{min}(K) \leq k(J) = 2 - \chi(J) = 2 - \chi(L) = 2 - 2 + 2g(L) = 2g_{min}(K) < 2g_{min}(K) + 1$, as required. ❑

12. ARBITRARILY LARGE MINIMUM GENUS

In Section 11, we found a general upper bound on the minimum crosscap number, based on the minimum genus. However, Auslander, Brown and Youngs in [1] have shown that for a graph G,

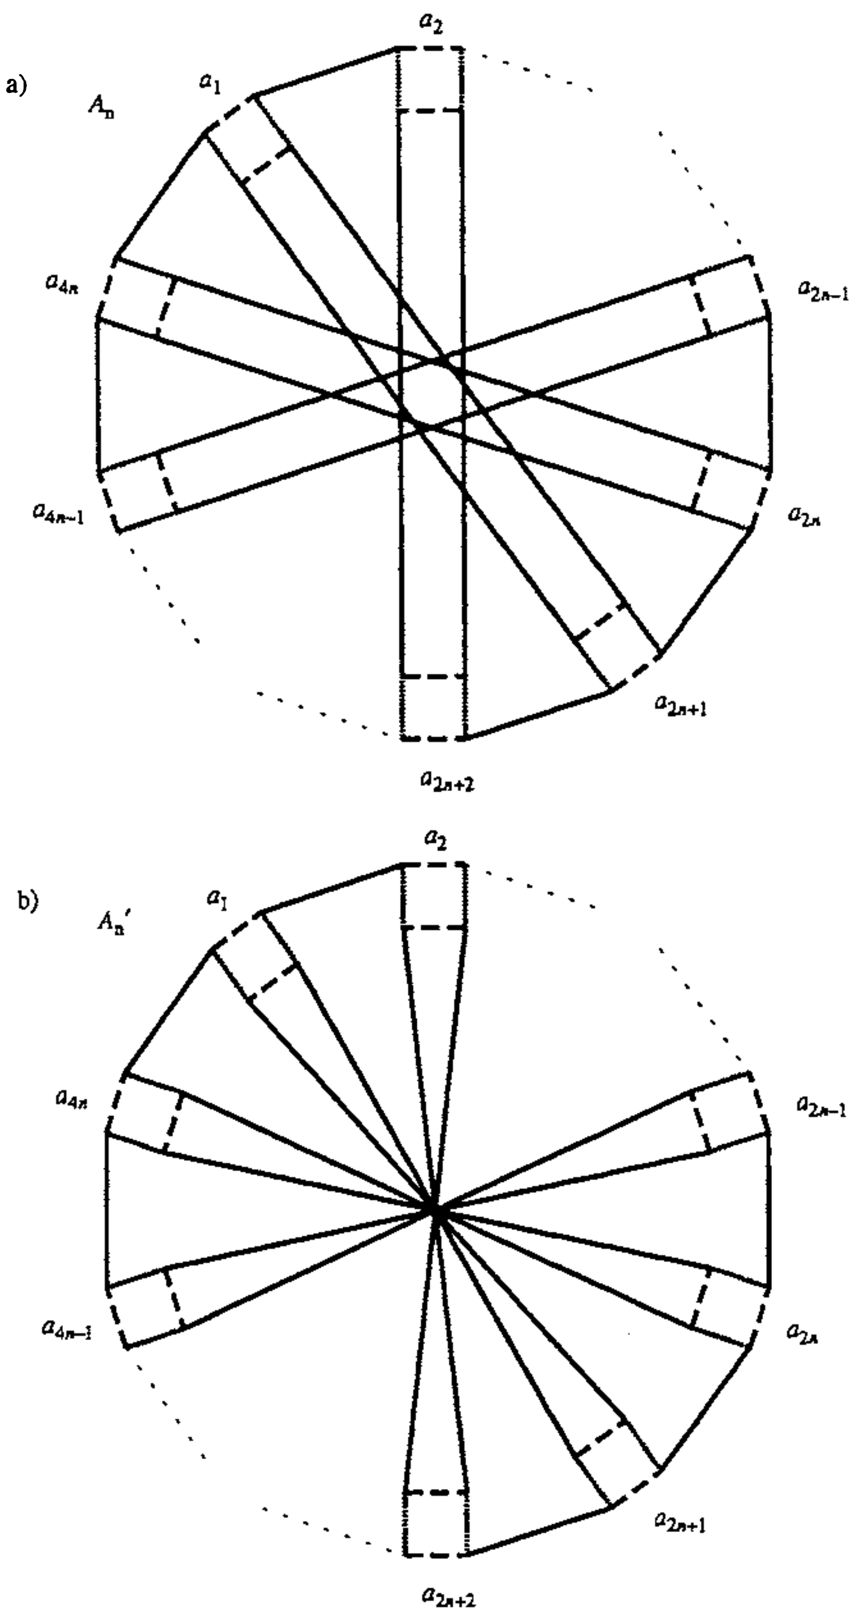
a)
A_n
a_2
a_1
a_{4n}
a_{2n-1}
a_{4n-1}
a_{2n}
a_{2n+1}
a_{2n+2}
b)
A_n'
a_2
a_1
a_{4n}
a_{2n-1}
a_{4n-1}
a_{2n}
a_{2n+1}
a_{2n+2}

FIGURE IV.11

Chapter V

MAXIMUM GENUS

1. INTRODUCTION

Calculating the minimum genus of a surface which admits a 2-cell imbedding of a given graph is still an open question. Recently, considerable attention has been paid to calculating the maximum genus of a surface which admits a 2-cell imbedding of a given graph G. (See [7-10, 15-19, 21, 22, 32, 33].) In our terminology, this amounts to determining the largest genus of a gem congruent to a gem K which G underlies. We will work in a more general setting and turn our attention to calculating, for a given 3-graph K, the highest possible genus $g_{max}(K)$ for a 3-graph congruent to K.

2. THE DEFICIENCY OF A GRAPH

The *deficiency* $\xi(G, T)$ of a spanning tree T for a connected graph G is defined to be the number of components of $G - ET$ that have an odd number of edges. Such components will be said to be *odd*. The *deficiency* $\xi(G)$ of the graph G is defined as the minimum of $\xi(G, T)$ over all spanning trees T.

THEOREM V.1 [9, 32]. *Let G be a connected graph. Then the minimum number of faces in any orientable imbedding of G is exactly $\xi(G) + 1$.* ❑

Theorem V.1 is often attributed to Xuong [32]. However, Khomenko had published it 6 years earlier. Our main theorem in this chapter is essentially a generalisation of Theorem V.1 to 3-graphs. This relationship between our theorem and Theorem V.1 is made evident in the final sections of this chapter.

3. THE DEFICIENCY OF A 3-GRAPH

Let A be a permitted red pole set in a 3-graph K such that $K - [A]$ has just one blue-yellow bigon. Then A is called a *pinch set* of K.

The *deficiency* $\xi(K, A)$ of a pinch set A for a connected 3-graph K is defined to be the number of red-yellow bigons in K that contain an odd number of edges in A. The *deficiency* $\xi(K)$ of the 3-graph K is defined as the minimum of $\xi(K, A)$ over all pinch sets A.

EXAMPLE V.2. *Consider the connected orientable* 3*-graph L in Figure* IV.1*a. The pinch sets of L are* $\{a_1\}$, $\{a_2\}$, $\{a_3\}$ *and* $\{a_4\}$. *In each case the deficiency of a pinch set is* 1, *and hence the deficiency of L is* 1. *There is no other orientable* 3*-graph congruent to L. Hence the minimum number of blue-yellow bigons in an orientable* 3*-graph congruent to L is* $2 = \xi(L) + 1$. *Theorem* V.14 *establishes the fact that this minimum is always* $\xi(L) + 1$ *for a general connected orientable* 3*-graph L.* ❑

The following theorem is the main theorem for this chapter and it will be proved in Section 6.

THEOREM V.14. *Let K be a connected orientable 3-graph. Then the minimum number of blue-yellow bigons in an orientable* 3*-graph congruent to K is* $\xi(K) + 1$.

4. SINGULAR 3-GRAPHS

An orientable 3-graph K with one blue-yellow bigon has genus $g_{max}(K)$. For this reason we give these 3-graphs a special name. An orientable 3-graph with just one blue-yellow bigon is said to be *singular*. It is immediate that a singular 3-graph is connected, since it has a circuit which passes through every vertex.

LEMMA V.3. *Let* $\{a, a'\}$ *be a permitted red pole set in an orientable* 3*-graph K. Suppose that a and* a' *belong to the same red-*

yellow bigon and that $K - [a, a']$ is singular. Then K is congruent to a singular 3*-graph.*

Proof. Consider the graph $K - [a]$. Label the vertices and edges of $K - [a]$ incident on or adjacent to a' as in Figure V.1a. Since $K - [a, a']$ has one blue-yellow bigon, then $K - [a]$ must have one or two blue-yellow bigons. Suppose that a' is a red 1-unipole in $K - [a]$. Therefore $\{c_1, b_1, c_2, b_2\} \subseteq R$ for some blue-yellow bigon R in $K - [a]$. If a' were consistent in $K - [a]$ then $K - [a, a']$ would have two blue-yellow bigons, a contradiction. Therefore a' must be inconsistent which, by Lemma IV.11, is also impossible since $K - [a]$ is orientable. We conclude that a' is a red 1-dipole in $K - [a]$ and that $K - [a]$ has two blue-yellow bigons. Let R_1 and R_2 denote the blue-yellow bigons in $K - [a]$ that include $\{c_1, b_1\}$ and $\{c_2, b_2\}$ respectively.

Label the vertices of K incident on neighbours of a as in Figure V.1a. Let b denote the edge in $K - [a]$ that joins the vertices v and w. (See Figure V.1b.) Without loss of generality, assume that $b \in R_1$. Since $K - [a]$ is bipartite, we may colour the vertices of $K - [a]$ black or white so that no two vertices of the same colour are adjacent. Assume without loss of generality that v is coloured black so that w is coloured white. Let L be the graph with vertex set $V(K - [a]) \cup \{v_2, w_2\}$ where $v_2, w_2 \notin V(K - [a])$, and edge set $(E(K - [a]) - \{b, c_2\}) \cup \{b_3, b_4, a_1, c_3, c_4\}$ such that $\psi a_1 = \{v_2, w_2\}$, $\psi b_3 = \{v, v_2\}$, $\psi b_4 = \{w, w_2\}$, $\psi c_3 = \{v_2, w_1\}$ and

$\psi c_4 = \{v_1, w_2\}$. (See Figure V.1c.) Hence L is obtained from K by reattaching a to c_2. Furthermore L is bipartite and contains only one blue-yellow bigon $R_3 = ((R_1 \cup R_2) - \{b, c_2\}) \cup \{c_3, b_3, c_4, b_4\}$. Hence L is the required orientable 3-graph. ❑

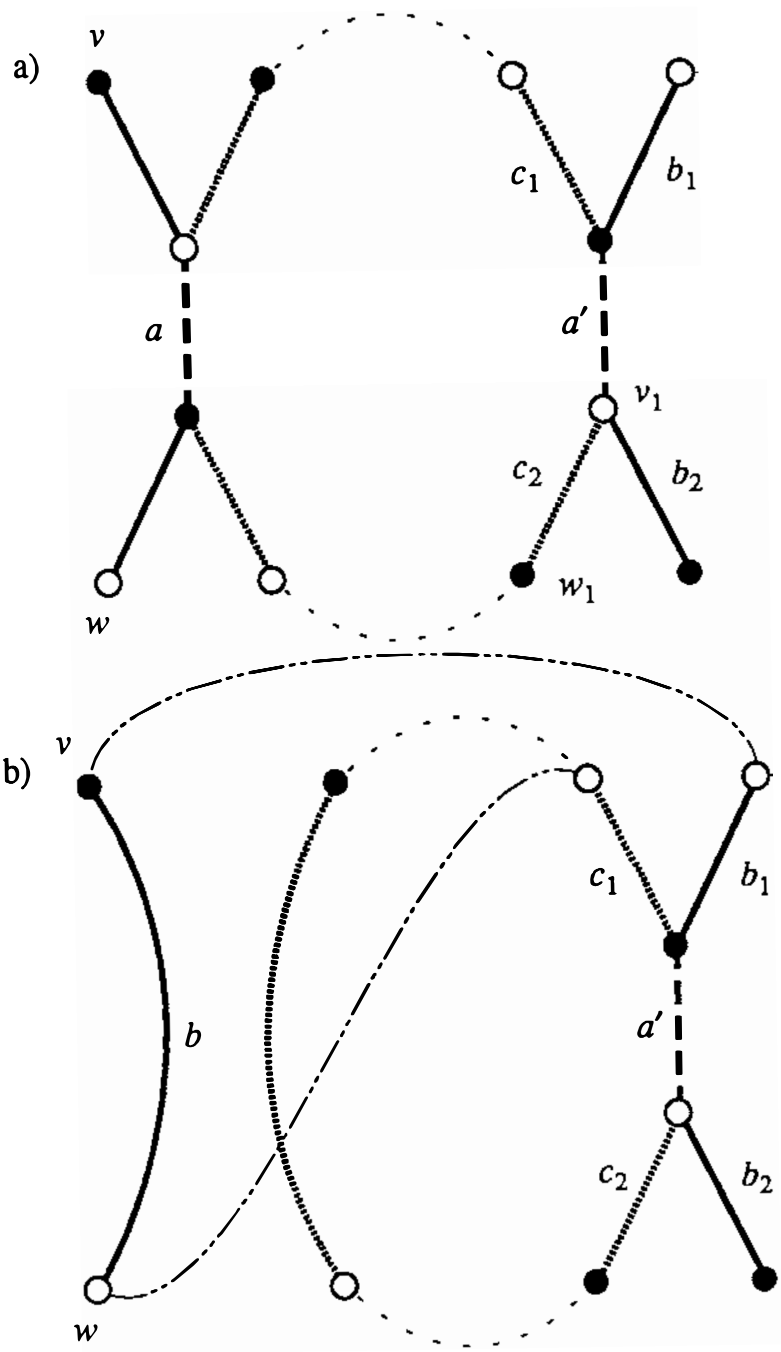

c)

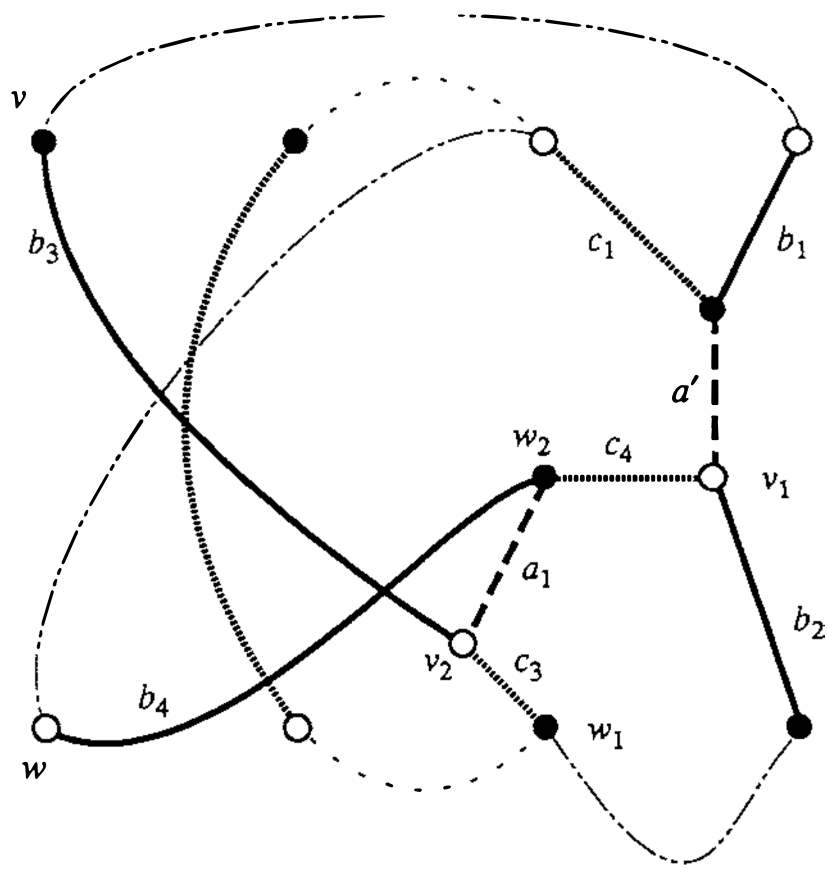

FIGURE V.1

Let K be a singular 3-graph and let R denote its one blue-yellow bigon. Let a and a' be distinct red edges in a red-yellow bigon B. Let a join v and w and let a' join v' and w'. If $v \in R_w[v', w']$ then we say that a and a' are *incoherent with respect to* R. (See Figure V.2.) A red-yellow bigon in K that contains a pair of incoherent edges is said to be *incoherent*. A red-yellow bigon that is not incoherent is *coherent*.

A singular 3-graph with an incoherent red-yellow bigon is said to be *weakly singular*. A singular 3-graph that lacks an incoherent red-yellow bigon is said to be *strongly singular*.

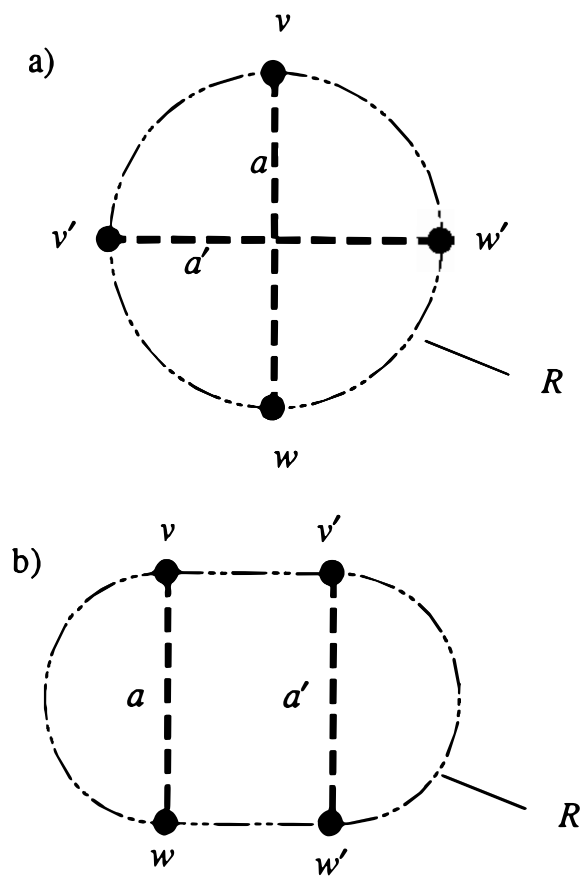

FIGURE V.2

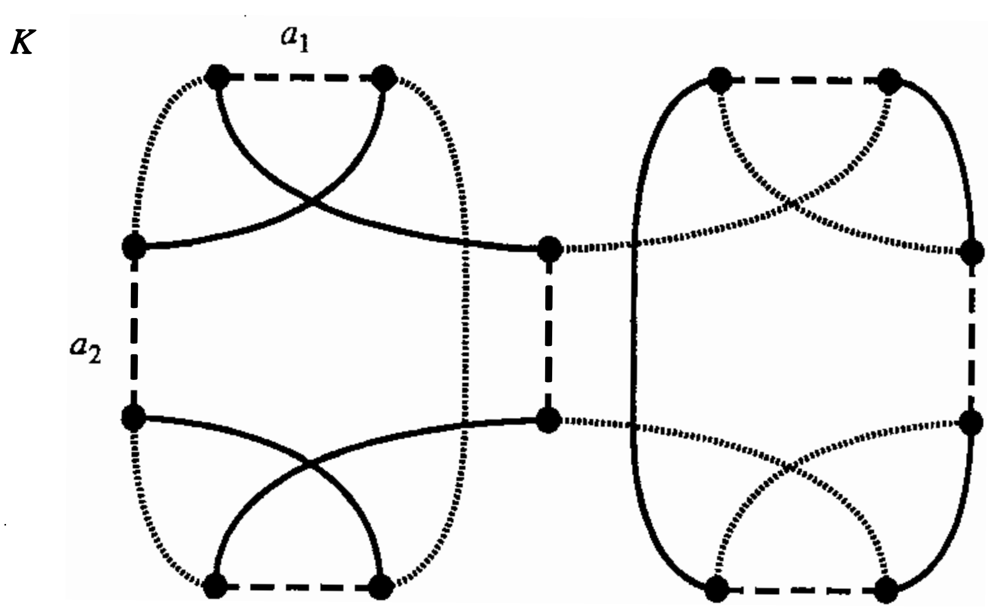

FIGURE V.3

EXAMPLE V.4. *Consider the singular* 3-*graph* K *in Figure* V.3. *The red edges* a_1 *and* a_2 *are incoherent, and hence* K *is weakly singular.* ❑

LEMMA V.5. *Let* K *be a weakly singular 3-graph with an incoherent red-yellow bigon* B. *Then there exists a permitted red pole set* $\{a, a'\} \subseteq B$ *such that* $K - [a, a']$ *is singular.*

Proof. Let R denote the blue-yellow bigon in K. Since B is incoherent, there exist red edges a and a' that are incoherent with respect to R. Clearly, if either a or a' belonged to a digon then the other edge would not be incoherent with it. Hence we conclude that a and a' are 1-poles in K. Since K is singular, then a and a' are 1-unipoles. Furthermore, by Lemma IV.11, a and a' are consistent 1-unipoles since singular 3-graphs are orientable. The following uses the notation of Figure V.4. Since a is consistent then $K - [a]$ has two blue-yellow bigons R_1 and R_2. Assume without loss of generality that $v' \in VR_1$. Since a and a' are incoherent in K, then $w' \in VR_2$ in $K - [a]$. Therefore a' is a 1-dipole in $K - [a]$ and hence R_1 and R_2 coalesce to form the one blue-yellow bigon of $K - [a, a']$, as required. ❑

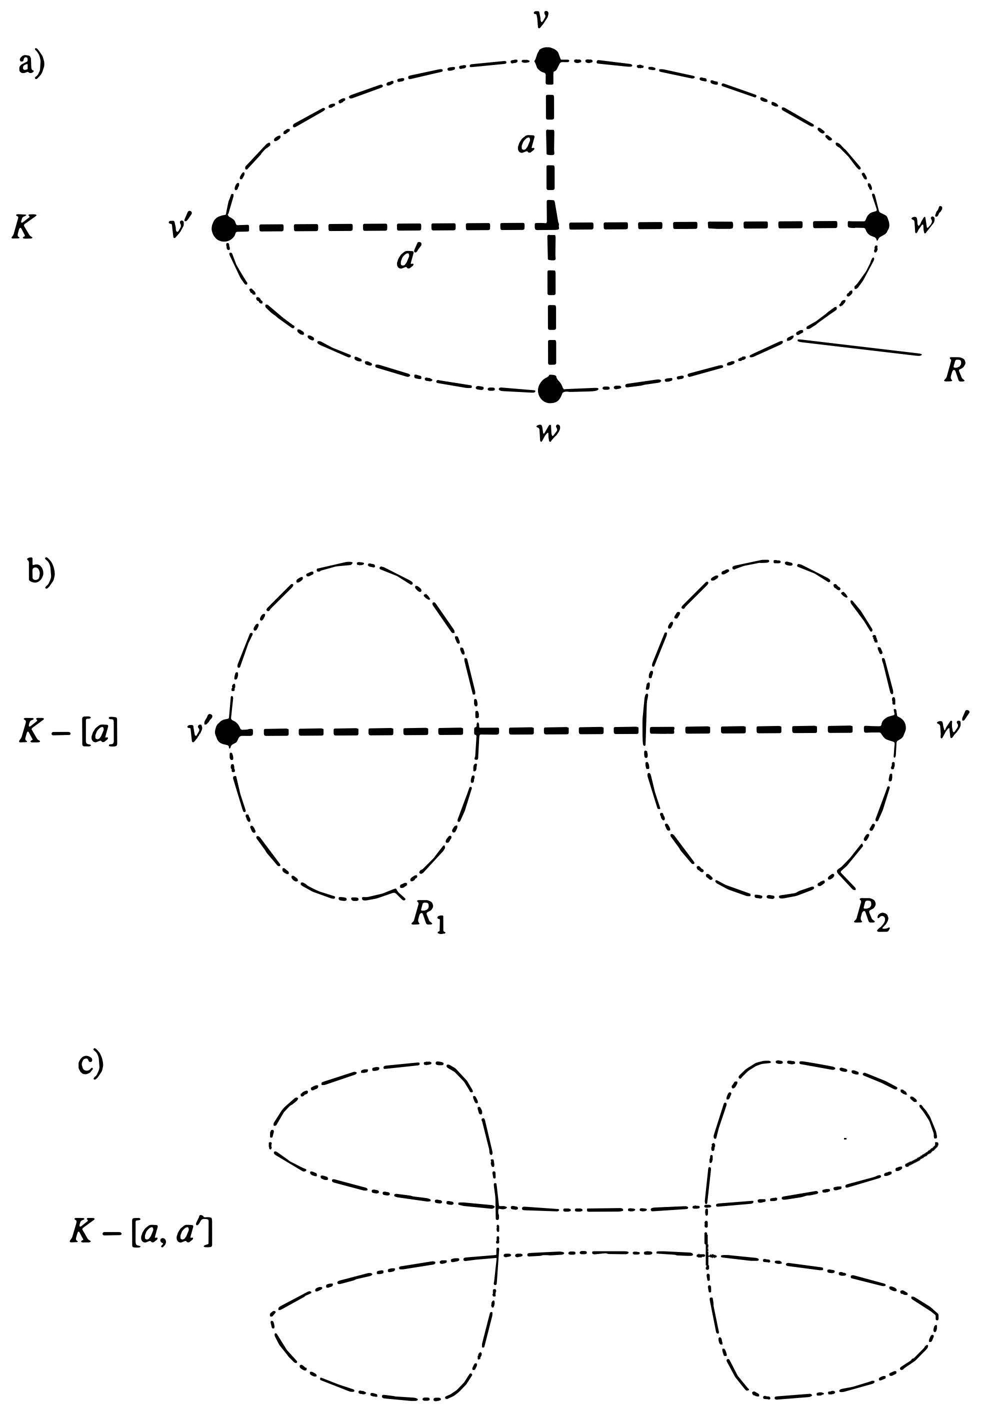

FIGURE V.4

EXAMPLE V.6. *Returning to Example* V.4, *we see that* $K - [a_1, a_2]$ *is singular, as expected.* (*See Figure* V.5.) *We note that* $K - [a_1, a_2]$ *is weakly singular since* a_3 *and* a_4 *are incoherent.* ❑

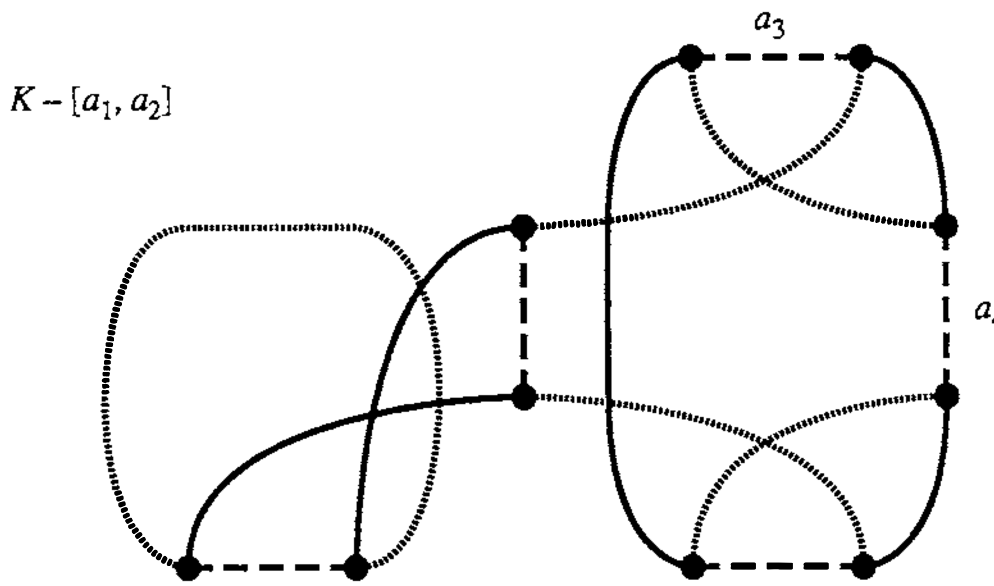

FIGURE V.5

LEMMA V.7. *All* 3-*graphs* ζ-*adjacent to a strongly singular* 3-*graph are singular.*

Proof. Let K denote a strongly singular 3-graph and let R denote the one blue-yellow bigon in K.

Firstly, assume that L is a 3-graph obtained from K by a reattachment of a red 1-pole a to a yellow edge c. The 1-pole a must be a 1-unipole since K is singular, and it must be consistent by Lemma IV.11. The following uses the notation of Figure IV.5. Let B denote the red-yellow bigon in $K - [a]$ that contains the yellow edge c'. Hence $c \in B$, by definition of cancellation. Since a is a consistent 1-unipole in K, let R_1 and R_2 denote the distinct blue-yellow bigons in $K - [a]$ that contain the edges b and c' respectively. Suppose $c \in R_1$. Since $c' \in R_2$ and $c, c' \in B$, we deduce that there exists a red edge $a_1 \in B$ that joins a vertex $v' \in VR_1$ to a vertex $w' \in VR_2$. Therefore $v \in R_w[v', w']$ and a and a_1 are incoherent, contradicting the fact that all red-yellow bigons in K are coherent.

Hence we conclude that $c \in R_2$. Now, since we have $b \in R_1$ and $c \in R_2$, then reattaching a to c therefore coalesces R_1 and R_2 into the one blue-yellow bigon $((R_1 \cup R_2) - \{b, c'\}) \cup \{b_1, c_1, b_2, c_2\}$ in L.

Now assume that L is obtained from K by a twist of a red cut edge a. The following uses the notation of Figure IV.5. Let R_1 and R_2 denote the blue-yellow bigons in $K - [a]$ that contain b and c' respectively. Evidently $(\{b_1, c_1, b_2, c_2\} \cup R_1 \cup R_2) - \{b, c'\}$ is the one blue-yellow bigon in L. ❑

5. ELEMENTARY 3-GRAPHS

An orientable 3-graph that is congruent to a strongly singular 3-graph but no weakly singular 3-graph is said to be *elementary*.

LEMMA V.8. *An elementary* 3-*graph is strongly singular.*

Proof. Let K be an elementary 3-graph. It suffices to show that it is singular, since by definition K cannot be weakly singular.

Let L be a strongly singular 3-graph congruent to K. By Theorem IV.33, K can be obtained from L by a finite sequence of ζ-moves. Let $L = L_1, L_2, ..., L_n = K$ be a sequence of 3-graphs such that L_{i+1} is ζ-adjacent to L_i whenever $1 \le i \le n - 1$. By Lemma V.7 L_2 is singular. However L_2 cannot be weakly singular, for otherwise the fact that K is elementary is contradicted. Hence L_2 is strongly singular and therefore, by Lemma V.7, L_3 is singular. Proceeding inductively we conclude that K is singular, as required. ❑

LEMMA V.9. *Any orientable 3-graph congruent to an elementary 3-graph is elementary.*

Proof. Let K be an orientable 3-graph congruent to an elementary 3-graph L. K is congruent to a strongly singular 3-graph since L is strongly singular by Lemma V.8. Now suppose K were congruent to a weakly singular 3-graph J. Then evidently L would be congruent to J also, contradicting the fact that L is elementary. Hence we conclude that K is elementary. ❑

LEMMA V.10. *Let K be an orientable 3-graph. Then there exists a pinch set A in K such that $K - [A]$ is elementary and $\xi(K) = \xi(K, A)$.*

Proof. Let A_1 be a pinch set in K such that $\xi(K, A_1) = \xi(K)$. Suppose $K - [A_1]$ is not elementary. Then there exists a weakly singular 3-graph L_1 congruent to $K - [A_1]$. Let B_1 denote an incoherent red-yellow bigon in L_1. By Lemma V.5, there exist red edges a_1 and a_1' in B_1 such that $L_1 - [a_1, a_1']$ is singular. Clearly $A_2 = A_1 \cup \{a_1, a_1'\}$ is a permitted red pole set in K. Furthermore $|B \cap A_2|$ is even if and only if $|B \cap A_1|$ is even, for any red-yellow bigon B in K. Hence $\xi(K, A_2) = \xi(K)$.

Now, if $K - [A_2]$ is not elementary then repeat the above proceedure with A_2. Proceeding inductively we eventually obtain a pinch set A_n such that $K - [A_n]$ is an elementary 3-graph and $\xi(K, A_n) = \xi(K)$. ❑

A pinch set in a 3-graph K such that $K - [A]$ is elementary is called an *elementary pinch set.*

$K - [a_1, a_2, a_3, a_4]$

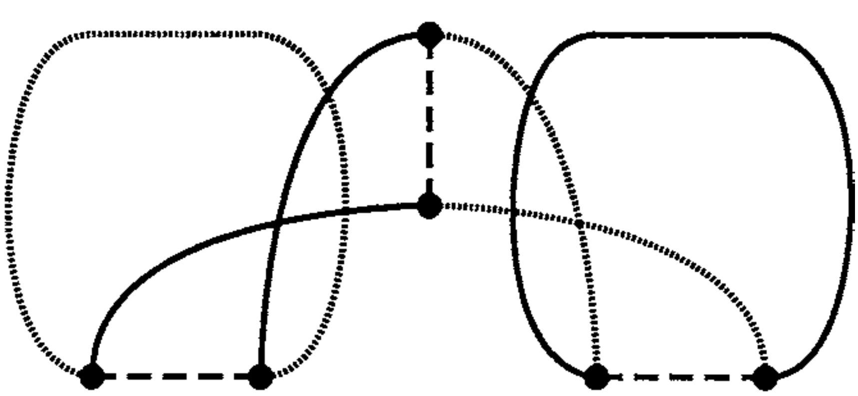

FIGURE V.6

EXAMPLE V.11. *Returning to Example* V.6, *we note that* $A = \{a_1, a_2, a_3, a_4\}$ *is a elementary pinch set of* K *since* $K - [A]$ *is strongly singular and is not congruent to a weakly singular* 3-*graph.* (*See Figure* V.6.) ❑

LEMMA V.12. *Let* K *and* L *be two congruent orientable* 3-*graphs. Then* $\xi(K) = \xi(L)$.

Proof. Let A be an elementary pinch set in K. By applying Lemma IV.6 recursively, we see that $K - [A]$ and $L - [A]$ are congruent orientable 3-graphs. This implies that $L - [A]$ is elementary by Lemma V.9. Therefore, by Lemma V.8, $L - [A]$ has just one blue-yellow bigon, and we conclude that A is a pinch set of L. Let B be a red-yellow bigon in K, and let B' be the red-yellow bigon in L that corresponds to B. Clearly $|B \cap A| = |B' \cap A|$ and hence $\xi(K, A) = \xi(L, A)$. Therefore $\xi(L) \leq \xi(L, A) = \xi(K, A) = \xi(K)$. A similar argument can be used to show that $\xi(K) \leq \xi(L)$. ❑

6. KHOMENKO'S THEOREM FOR 3-GRAPHS

LEMMA V.13. *A connected orientable* 3-*graph* K *is congruent to a singular* 3-*graph if and only if* $\xi(K) = 0$.

Proof. Firstly, assume that K is congruent to a singular 3-graph K'. Consider the following cases.

i) K' is elementary. By Lemma V.9, K is elementary and so by Lemma V.8 it is singular. Therefore $\varnothing$ is a pinch set of K and hence $\xi(K) = 0$.

ii) K' is not elementary. Then there exists a weakly singular 3-graph K'' congruent to K with an incoherent red-yellow bigon B. By Lemma V.5 there exist red edges a and a' in B such that $K'' - [a, a']$ is singular. Hence $\{a, a'\}$ is a pinch set in K'', and since $|\{a, a'\} \cap B|$ is even, we conclude that $\xi(K'') = \xi(K'', \{a, a'\}) = 0$. By Lemma V.12, $\xi(K) = \xi(K'') = 0$, as required.

Now assume that $\xi(K) = 0$. Let $A_1 = \{a_1, a_2, ..., a_n\}$ be a pinch set such that $\xi(K, A_1) = 0$. Since $|B \cap A_1|$ is even for all red-yellow bigons B in K, n is even, and we may label the elements of A_1 so that a_i and a_{i+1} belong to the same red-yellow bigon for all odd i. Consider $K - [a_3, a_4, ..., a_n]$. Since $K - [A_1]$ is singular and a_1 and a_2 belong to the same red-yellow bigon, then by Lemma V.3, $K - [a_3, a_4, ..., a_n]$ is congruent to a singular 3-graph L_1. Consider $K - [a_5, a_6, ..., a_n]$. Since L_1 is singular, and a_3 and a_4 belong to the same red-yellow of L_1, then by Lemma V.3, $L_1 - [a_3, a_4]$ is congruent to a singular 3-graph L_2. Furthermore, $L_1 - [a_3, a_4]$ is congruent to $K - [a_3, a_4, a_5, a_6, ..., a_n] - [a_3, a_4] = K - [a_5, a_6, ..., a_n]$, and hence $K - [a_5, a_6, ..., a_n]$ is congruent to the singular 3-graph L_2.

Proceeding inductively, the fact that n is even implies that K is congruent to a singular 3-graph, as required. ❑

THEOREM V.14. *Let K be a connected orientable 3-graph. Then the minimum number of blue-yellow bigons in an orientable 3-graph congruent to K is $\xi(K) + 1$.*

Proof. It suffices to prove that K is congruent to an orientable 3-graph with $n + 1$ or fewer blue-yellow bigons if and only if $\xi(K) \leq n$. We prove this by induction on n. By Lemma V.13 the theorem holds for $n = 0$, so we assume that it holds for all $k < n$, where $n > 0$.

Firstly, suppose L is an orientable 3-graph congruent to K that has $n + 1$ blue-yellow bigons. Hence there exists a red 1-dipole a in L, for otherwise L would have just one blue-yellow bigon, or more than one component. Therefore $L - [a]$ has n blue-yellow bigons. By the induction hypothesis, $\xi(L - [a]) \leq n - 1$. Choose a pinch set A of $L - [a]$ such that $\xi(L - [a], A) = \xi(L - [a])$. Clearly $A \cup \{a\}$ is a pinch set of L. Let B be the red-yellow bigon in L containing a and let B' be the red-yellow bigon in $L - [a]$ that corresponds to B. Consider the following two cases.

a) If $|B' \cap A|$ is odd then $|B \cap (A \cup \{a\})|$ is even and hence $\xi(L, A \cup \{a\}) = \xi(L - [a], A) - 1 = \xi(L - [a]) - 1$. Therefore $\xi(L) \leq \xi(L, A \cup \{a\}) = \xi(L - [a]) - 1 \leq n - 2 < n$.

b) If $|B' \cap A|$ is even then $|B \cap (A \cup \{a\})|$ is odd and hence $\xi(L, A \cup \{a\}) = \xi(L - [a], A) + 1 = \xi(L - [a]) + 1$. Therefore $\xi(L) \leq \xi(L, A \cup \{a\}) = \xi(L - [a]) + 1 \leq n$.

Hence by Lemma V.12, $\xi(K) = \xi(L) \leq n$, as required.

Now suppose that $\xi(K) = n$. Let A be a pinch set of K such that $\xi(K) = \xi(K, A) = n$. Let $a \in A$ be a red 1-pole in a red-yellow bigon B, where $|B \cap A|$ is odd. (B exists since $n > 0$.) Let B' denote the red-yellow bigon in $K - [a]$ that corresponds to B. Clearly $A - \{a\}$ is a pinch set of $K - [a]$. It is also clear that $\xi(K - [a]) \leq \xi(K - [a], A - \{a\}) = \xi(K, A) - 1 = n - 1$, since $|B' \cap (A - \{a\})|$ is even. By the induction hypothesis, $K - [a]$ is congruent to an orientable 3-graph L which has no more than n blue-yellow bigons. It is immediate that K is congruent to an orientable 3-graph which has no more than $n + 1$ blue-yellow bigons, as required. ❑

7. ELEMENTARY GEMS

In the next three sections we show how elementary gems correspond to imbeddings of trees in surfaces. (Recall that a tree is a connected graph in which every edge is an isthmus.) For the present section we suppose K to be a gem and G the graph which underlies K.

Let Y be a red-blue bisquare in K, with red edges a and a'. Furthermore, suppose a and a' are 1-poles in K. Evidently, a' belongs to a red-blue 2-dipole $\{a', b\}$ in $K - [a]$, and therefore $K' = K - [a] - [a', b]$ is a gem. Let e be the edge in G that corresponds to Y. We observe that $G - \{e\}$ underlies K'. We say that e *corresponds* to both a and a'. Moreover, we say that a and a' both *represent* e.

LEMMA V.15. *Let G be the graph that underlies a gem K, and suppose a is a red cut edge in K. Then the edge e in G that corresponds to a is an isthmus.*

Proof. Let Y be the red-blue bisquare in K that contains a, and let a' be the red edge in Y other than a. Consider the following cases.

i) a' is not a red 1-pole in K. Then clearly a' belongs to a red-yellow 2-dipole in K. This implies that e is incident on a vertex which has degree one and hence e is an isthmus.

ii) a' is a red 1-pole in K. Let $\{a', b\}$ be the red-blue 2-dipole in $K - [a]$ that contains a'. Then $G - \{e\}$ is the graph that underlies $K - [a] - [a', b]$. Hence $c(G - \{e\}) = c(K - [a] - [a', b]) > c(K) = c(G)$, since a is a red cut edge in K. This implies that e is an isthmus in G. ❑

EXAMPLE V.16. *Consider the graph G in Figure* V.7*a that underlies the gem K in Figure* V.7*b. Since each red-blue bisquare in K contains a red cut edge, then by Lemma* V.15, *G is a tree.* ❑

a)

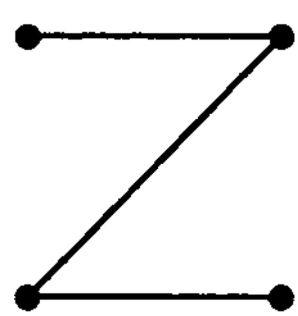

b)

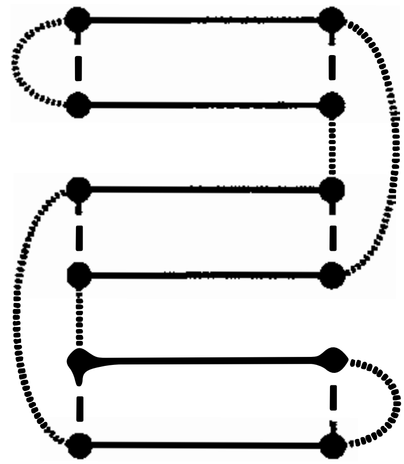

FIGURE V.7

LEMMA V.17. *Let G be the graph that underlies a gem K, and suppose e is an isthmus in G. Let a be a red edge in K that represents e. Then either a is a red cut edge or a belongs to a red-yellow 2-dipole in K.*

Proof. This proof is similar to the proof of Lemma V.15, and is omitted. ❑

8. CAPS AND CROSSCAPS OF BLUE-YELLOW BIGONS

Let K be a gem with just one blue-yellow bigon R. Let Y be an arbitrary red-blue bigon of K. Label the edges and vertices of Y as in Figure V.8a. If $w' \in V(R_v[v', w])$ then we say Y is a *cap* (of R); otherwise Y is a *crosscap* (of R). If K is orientable then clearly Y is a cap, for otherwise a would belong to a circuit $R_{w'}[v, w] \cup \{a\}$ of odd length. Assume Y is a cap, and suppose there exists another cap X

of R such that $|X \cap R_v[v', w]| = 1$. Then we say that X and Y are *bound* in R, and $\{X, Y\}$ is a *clamp* of R. (See Figure V.8b.)

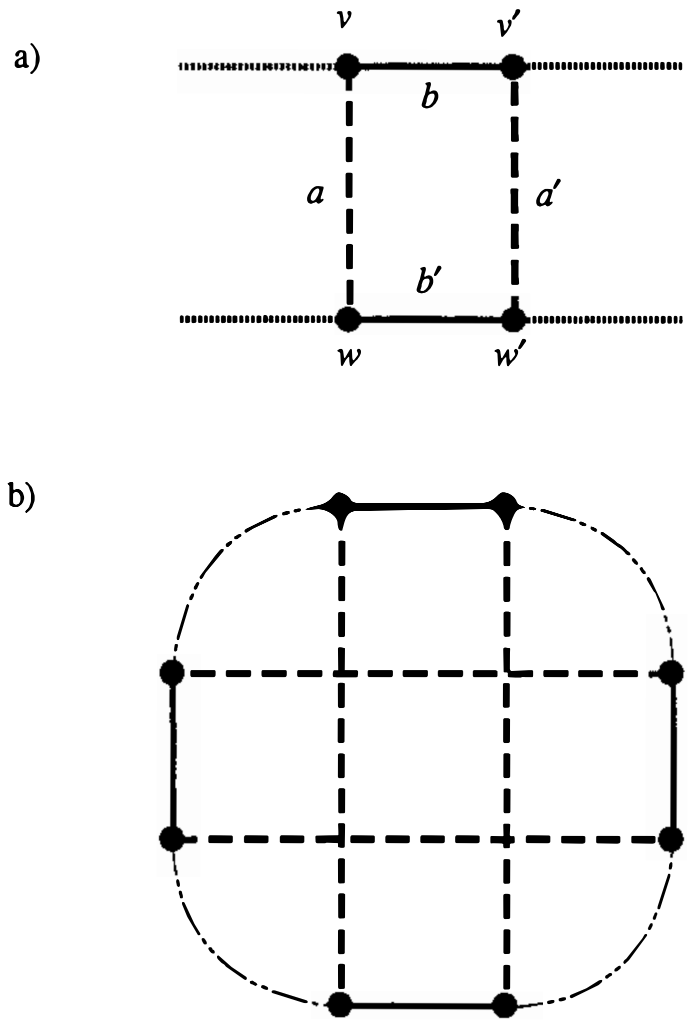

FIGURE V.8

LEMMA V.18. *Let K be a singular gem with blue-yellow bigon R. If there exists a red 1-pole a in K that is not a red cut edge, then there exists a clamp of R.*

Proof. Let a join v and w, and let P_1 and P_2 be the two disjoint paths in R that join v and w. The following uses the notation of Figure V.9. Since $c(K - [a]) = 1$, then there exists a red edge a_1 other than a that joins a vertex $v_1 \in VP_1$ to a vertex $w_1 \in VP_2$. Clearly a and a_1 belong to distinct caps of R, for otherwise they would belong to a crosscap of R, contradicting the fact that K is orientable. Let Y be the cap of R that contains a and let X be the cap of R that contains a_1. Let b_1 be the blue edge of X incident on w_1. Since $v_1 \in VP_1$ and $w_1 \in VP_2$, then evidently $|X \cap R_v[v', w]| = |\{b_1\}| = 1$. Hence $\{X, Y\}$ is a clamp of R, as required. ❑

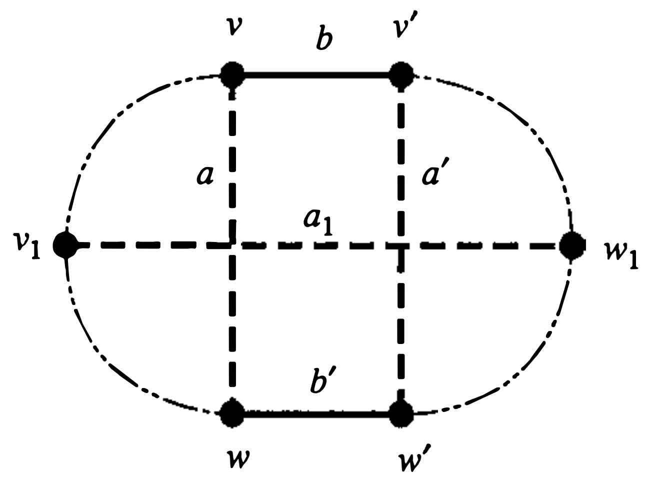

FIGURE V.9

EXAMPLE V.19. *Consider the singular gem K in Figure* V.10 *and let R denote the one blue-yellow bigon in K. Since a is not a red cut edge, then by Lemma* V.18, *a set comprising Y and some other red-blue bisquare is a clamp of R. We note that $v_1 \in VR_{v'}[v, w]$ and*

$w_1 \in VR_{v'}[v, w]$, and hence $\{Y, Y_1\}$ is a clamp of R, where Y_1 is the red-blue bisquare that contains a_1. ❑

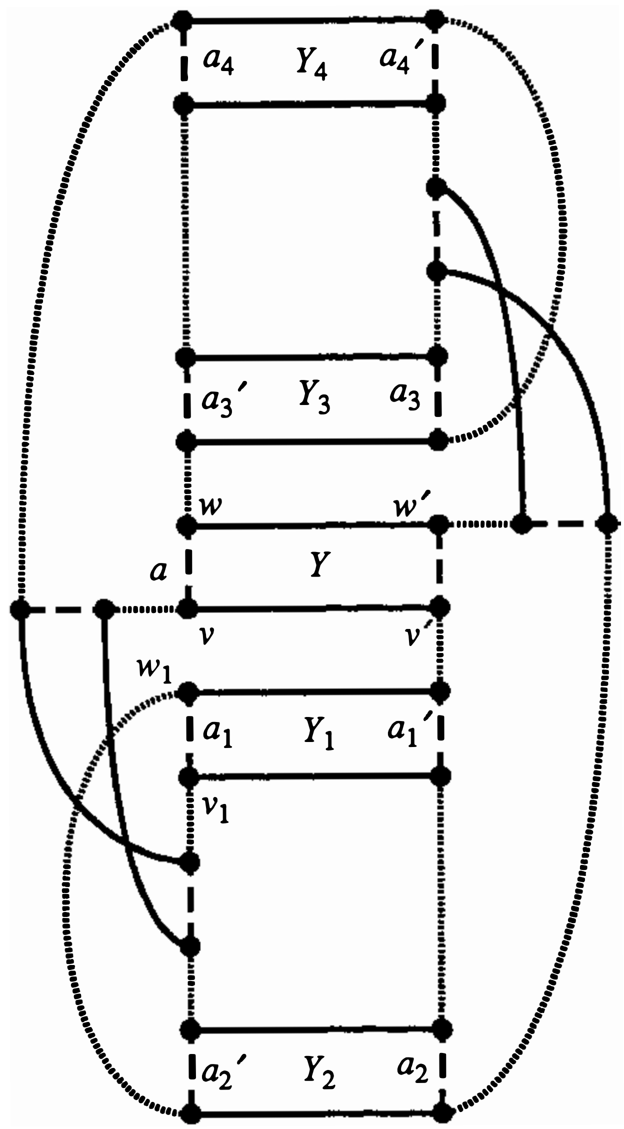

FIGURE V.10

LEMMA V.20. *If K is an elementary gem then every red 1-pole in K is a red cut edge.*

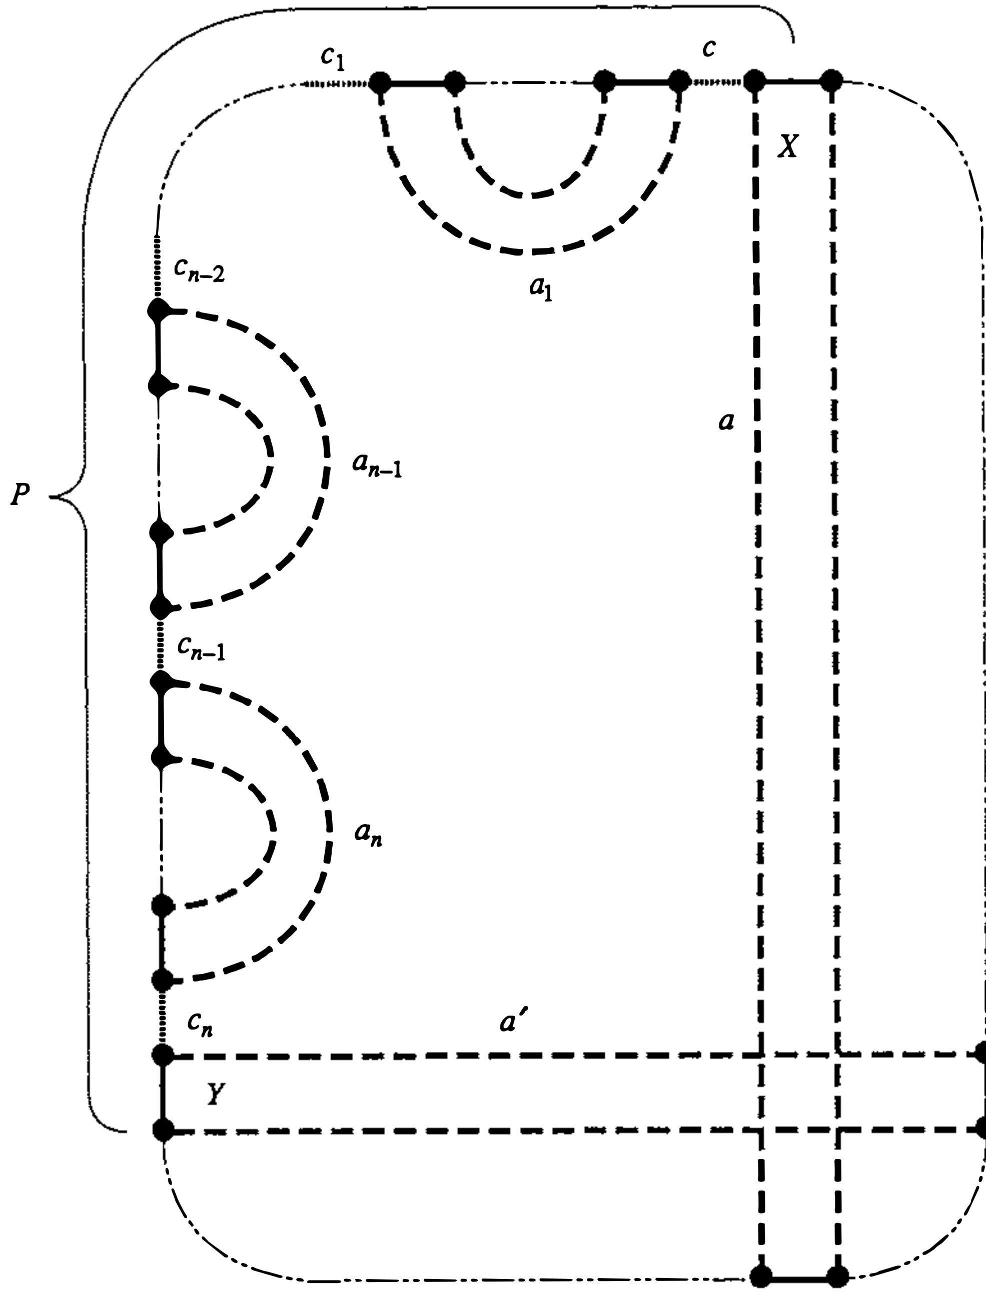

FIGURE V.11

Proof. Let R be the one blue-yellow bigon in K and suppose there is a red edge in K that is not a red cut edge. By Lemma V.18, there exists a clamp $\{X, Y\}$ of R in K. Choose X and Y so that the length of a path P in R, with one terminal edge in X and the other in

Y, is minimised. We note that no clamp $\{X', Y'\}$ exists such that $|(X' \cup Y') \cap P| > 1$, for otherwise we contradict the minimality of $|P|$. The following uses the notation of Figure V.11. Clearly $\{a, c, a_1, c_1, \ldots, c_{n-1}, a_n, c_n, a'\}$ is a path of red and yellow edges, and hence a and a′ are incoherent, contradicting the fact that K is elementary. ❑

COROLLARY V.21. *Let G be the graph that underlies a gem K. Then G is a tree if and only if K is an elementary gem.* ❑

9. SEMI-GEMS

A *semi-gem* is a 3-graph such that each red-blue bigon is a circuit of length 2 or 4. Hence in a semi-gem, the red-blue bigons that are not bisquares are red-blue 2-dipoles. By Lemma IV.5, any 3-graph congruent to a semi-gem is a semi-gem. Furthermore, if A is a permitted red pole set in a gem K, then $K - [A]$ is a semi-gem.

EXAMPLE V.22. *Consider the pinch set* $A = \{a_1, a_2, a_3, a_4\}$ *of the gem K in Figure* V.10. *Then the semi-gem in Figure* V.12 *is* $K - [A]$. ❑

Cancelling all the red-blue 2-dipoles in a semi-gem K results in a gem, called the *frame* of K. A graph *underlies* a semi-gem K if it underlies the frame of K.

We generalise Corollary V.21 with the following lemma. Its proof is immediate.

LEMMA V.23. *Let G be the graph that underlies a semi-gem K. Then G is a tree if and only if K is an elementary semi-gem.* ❑

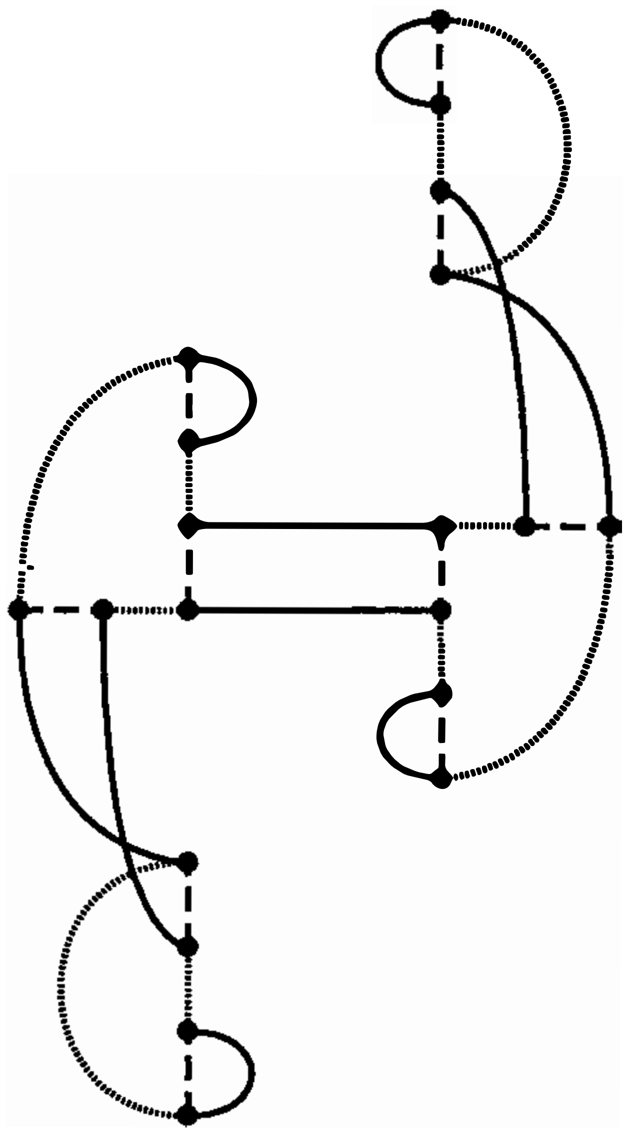

FIGURE V.12

EXAMPLE V.24. *Consider the semi-gem $K - [A]$ in Figure* V.12. *Then the gem in Figure* V.7*b is the frame of $K - [A]$ and the graph G in Figure* V.7*a underlies $K - [A]$. Since $K - [A]$ is an* *elementary semi-gem, then by Lemma* V.23 *G is a tree, as expected.* ❑

Suppose $A = \{a_1, a_2, ..., a_n\}$ is a permitted pole set in a gem K. Then $|Y \cap A| = 0$ or 1 for any red-blue bisquare Y in K. Let $Y = \{Y_1, Y_2, ..., Y_n\}$ be the set of red-blue bigons where $Y_i \cap A = \{a_i\}$ whenever $1 \leq i \leq n$. We say that Y *corresponds* to A. Furthermore, let $E = \{e_1, e_2, ..., e_n\}$ be the set of edges in the graph that underlies K, where e_i corresponds to a_i whenever $1 \leq i \leq n$. We say that E *corresponds* to A and to Y. If a_k' is the red edge in Y_k other than a_k for some k, then clearly both E and Y correspond to $A + \{a_k, a_k'\} = A + \rho(Y_k)$.

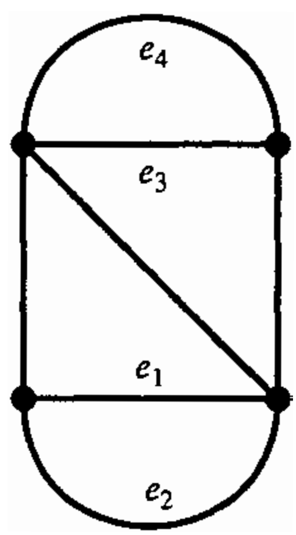

FIGURE V.13

EXAMPLE V.25. *Consider the gem K in Figure* V.10, *and let A be the pinch set* $\{a_1, a_2, a_3, a_4\}$. *Clearly the graph G in Figure* V.13, *underlies K. Furthermore* $Y = \{Y_1, Y_2, Y_3, Y_4\}$ *and* $E = \{e_1, e_2, e_3, e_4\}$ *both correspond to A.* ❑

LEMMA V.26. *Suppose A is a pinch set in a gem K and Y is a red-blue bigon that meets A. Then $A + \rho(Y)$ is a pinch set of K.*

Proof. Label the edges of Y as in Figure V.8a and assume without loss of generality that $a \in A$. Hence $a' \notin A$. Consider $K' = K - [(A - \{a\})]$. Clearly $K' - [a']$ has just one blue-yellow bigon since $K' - [a]$ has just one. Hence $\{a'\}$ is a pinch set of K', which implies $(A - \{a\}) \cup \{a'\} = A + \rho(Y)$ is a pinch set of K, as required. ❑

EXAMPLE V.27. *Let A be the pinch set $\{a_1, a_2, a_3, a_4\}$ in the gem K of Figure* V.10. *Then by Lemma* V.26 $A + \{a_k, a_k'\}$ *is also a pinch set of K, whenever* $1 \leq k \leq 4$. ❑

10. THE PRINCIPAL PARTITION

Let Y be the set of red-blue bigons that correspond to a permitted red pole set A in a gem K. Let E be the set of edges that correspond to A in the graph G that underlies K. Let $G_1, G_2, ..., G_m$ be the components of $G[E]$. Then we have a partition $\mathbb{Y} = \{Y_1, Y_2, ..., Y_m\}$ of Y where the red-blue bigons in Y_i corresponds to the edges in EG_i whenever $0 \leq i \leq m$. We say that $\mathbb{Y}$ is the *principal partition* of Y.

EXAMPLE V.28. *Returning to Example* V.25, *we have that* $\mathbb{Y} = \{\{Y_1, Y_2\}, \{Y_3, Y_4\}\}$ *is the principal partition of* Y *since* $G[\{e_1, e_2\}]$ *and* $G[\{e_3, e_4\}]$ *are the two components of* $G[E]$. ❑

LEMMA V.29. *The sets* $N(\bigcup Y_1)$, $N(\bigcup Y_2)$, ..., $N(\bigcup Y_m)$ *are pairwise disjoint.*

Proof. Suppose the lemma is false, and let B be a red-yellow bigon that belongs to $N(\bigcup Y_i)$ and $N(\bigcup Y_j)$ where $i \neq j$, $1 \leq i \leq m$ and $1 \leq j \leq m$. Evidently the vertex v in VG that corresponds to B is a vertex of both VG_i and VG_j. This contradicts the fact that G_i and G_j are distinct components of $G[E]$. ❑

COROLLARY V.30. *If* B *is a red-yellow bigon that meets* A, *then* $B \cap A \subseteq \bigcup Y_i$ *for exactly one value of* i. ❑

For each i, we define $t_i(A)$ to be the number of red-yellow bigons in $N(\bigcup Y_i)$ that meet A in an odd number of edges. The following result follows directly from Corollary V.30.

EXAMPLE V.31. *Returning to Example* V.28, *we have that* $t_1(A) = 2$, *and* $t_2(A) = 2$. ❑

COROLLARY V.32. *If* A *is a pinch set, then* $\xi(K, A) = \sum_i t_i(A)$. ❑

LEMMA V.33. $t_i(A) \equiv |Y_i|$ *(mod 2) whenever* $1 \leq i \leq m$.

Proof. This lemma follows from the fact that

$$|Y_i| = \sum_{B \in N(\bigcup Y_i)} |B \cap A|. \ \square$$

11. RELATING THE DEFICIENCES OF GEMS AND GRAPHS

We conclude this chapter with the relationship between the deficiency of a gem and the underlying graph. We employ the notation of the previous section.

LEMMA V.34. *If A is a pinch set such that $\xi(K, A) = \xi(K)$, then $t_i(A) = 0$ or 1 whenever $0 \le i \le m$.*

Proof. Suppose $t_i(A) \ge 2$ for some value of i. Let B_1 and B_2 be two red-yellow bigons in $N(\bigcup Y_i)$ such that $|B_1 \cap A|$ and $|B_2 \cap A|$ are both odd. Since G_i is connected, let P be a path in G_i from the vertex v that corresponds to B_1 to the vertex w that corresponds to B_2. Let $Y_1, Y_2, \ldots, Y_{|P|}$ be the red-blue bigons that correspond to the edges in P. Each of these meets A, since $P \subseteq EG$. Let

$$A' = A + \sum_{i=1}^{|P|} \rho(Y_i).$$

By applying Lemma V.26 inductively, we see that A' is a pinch set of K. Furthermore $|B_1 \cap A'|$ and $|B_2 \cap A'|$ are both even, and $|B \cap A'| \equiv |B \cap A| \pmod 2$ for all $B \in B(K) - \{B_1, B_2\}$. Hence $\xi(K, A') = \xi(K, A) - 2 = \xi(K) - 2$, a contradiction. $\square$

EXAMPLE V.35. *Returning to Example* V.31, *the fact that* $t_1(A) = t_2(A) = 2$ *implies that* $\xi(K, A) > \xi(K)$ *by Lemma* V.34. *From the proof of Lemma* V.34, *we can construct a pinch set* A' *from* A *such that* $\xi(K, A') = \xi(A)$. *This can be done by adding to* A *the set* $\{a_1, a_1', a_3, a_3'\}$. ❑

LEMMA V.36. *Let* K *be an orientable connected gem, and let* G *be the graph that underlies* K. *Then* $\xi(G) \leq \xi(K)$.

Proof. Let A be an elementary pinch set in K such that $\xi(K) = \xi(K, A)$. Let Y, E, $G_1, G_2, ..., G_m$, and $\mathbb{Y}$ be defined as in Section 10. Evidently $K - [A]$ is an elementary semi-gem, and by Lemma V.23, $G - E$ is a spanning tree T. By definition, $\xi(G, T)$ is the number of odd components in $\{G_1, G_2, ..., G_m\}$, and so it is the number of cells in $\mathbb{Y}$ with odd cardinality. By Lemmas V.33 and V.34 and Corollary V.32 we have $\xi(G) \leq \xi(G, T) = \sum_i t_i(A) = \xi(K, A) = \xi(K)$, as required. ❑

LEMMA V.37. *Let* G *be the graph that underlies a connected gem* K. *Then* $\xi(G) \geq \xi(K)$.

Proof. Let T be a spanning tree in G such that $\xi(G, T) = \xi(G)$. Let $Y = \{Y_1, Y_2, ..., Y_m\}$ be the set of red-blue bigons in K that correspond to the edges in $EG - ET$. Let $\mathbb{Y} = \{Y_1, Y_2, ..., Y_n\}$ be the principal partition of Y. Let a_i be a red edge in Y_i for each i, and let $A = \{a_1, a_2, ..., a_m\}$. Then A is a pinch set, since T is a spanning tree. From the proof of Lemma V.34, we may choose A so that $t_i(A) = 0$ or 1 for each i. With such a set A, $t_i(A) = 1$ if and only if the component of $G - T$ corresponding to Y_i is odd. Hence $\xi(K) \leq \xi(K, A) = \sum_i t_i(A) = \xi(G, T) = \xi(G)$, as required. ❑

THEOREM V.38. *Let G be the graph that underlies a gem K. Then $\xi(G) = \xi(K)$.*

Proof. This theorem follows directly from Lemmas V.36 and V.37. ❑

Chapter VI

IRREDUCIBLE DOUBLE COVERED GRAPHS

1. INTRODUCTION

In [30], short proofs of three graph theoretic versions of the Jordan curve theorem are given. Our main theorem in this chapter, Theorem VI.7, generalises the version expressed in terms of a double cover for a graph. Two combinatorial generalisations of the Jordan curve theorem are shown to be equivalent in [14]. Moreover, one version is in terms of 3-graphs. In the same spirit, we discuss the equivalence of special cases of Theorems III.15 and VI.7.

2. DOUBLE COVERS

A family $\boldsymbol{D} = (D_1, D_2, ..., D_n)$ of cycles in a non-empty graph G is said to be a *cycle double cover* if for every edge e in G there exist exactly two values of i for which $e \in D_i$. Furthermore, $\boldsymbol{D}$ is a *circuit double cover* if each D_i is a circuit in G and we say that the ordered pair $(G, \boldsymbol{D})$ is a *(circuit) double covered graph.* Henceforth we take the term "double cover" to mean a circuit double cover, unless an indication to the contrary is given. Moreover properties of G, such as connectedness, may also be ascribed to $(G, \boldsymbol{D})$.

EXAMPLE VI.1. *Let $(K, \boldsymbol{P}, \boldsymbol{O})$ be a 3-graph where $\boldsymbol{P}$ is a proper edge colouring in three colours of a cubic graph K and $\boldsymbol{O}$ is an ordering of the colours. Let $\boldsymbol{D}$ denote the set of bigons in $(K, \boldsymbol{P}, \boldsymbol{O})$. Since each edge in K appears in just two distinct bigons, $(K, \boldsymbol{D})$ is a circuit double covered graph.* ❑

LEMMA VI.2. *If $(G, \boldsymbol{D})$ is a double covered graph then $\sum \boldsymbol{D} = \varnothing$.*

Proof. Since each edge in EG appears twice in the sum $\sum \boldsymbol{D}$ then it is immediate that $\sum \boldsymbol{D} = \varnothing$. ❑

Let $(G, \boldsymbol{D})$ be a double covered graph. If J is a subgraph of G, let $\boldsymbol{D}_J$ denote the subset of $\boldsymbol{D}$ consisting of those circuits contained in J. The pair $(G, \boldsymbol{D})$ is called a *reducible* double covered graph of G if there exists a non-empty subgraph J of G such that $(J, \boldsymbol{D}_J)$ is a double covered graph. The pair $(G, \mathbf{D})$ is an *irreducible* double covered graph if it is not reducible. If J is a non-empty subgraph of G such that $(J, \mathbf{D}_J)$ is irreducible, then J is an *irreducible component* of G (with respect to $\mathbf{D}$.)

If $(G, \mathbf{D})$ is irreducible then we define the *Euler characteristic* of $(G, \mathbf{D})$ to be $\chi(G, \mathbf{D}) = |\mathbf{D}| - |EG| + |VG|$.

EXAMPLE VI.3. *Consider the graph G of Figure* VI.1. *Let* $D_1 = \{a, b\}$ *and* $D_2 = \{c, d\}$ *and let* $\mathbf{D}$ *be the family* (D_1, D_2, D_1, D_2). *Evidently* $(G, \mathbf{D})$ *is a double covered graph. However, the fact that* $(G[\{a, b\}], (D_1, D_1))$ *is a double covered graph implies that* $(G, \mathbf{D})$ *is reducible. In fact,* $G[\{a, b\}]$ *and* $G[\{c, d\}]$ *are the two irreducible components of G with respect to* $\mathbf{D}$. *Now let* $D_3 = \{a, d\}$ *and* $D_4 = \{b, c\}$ *and let* $\mathbf{D}' = (D_1, D_2, D_3, D_4)$. *Again,* $(G, \mathbf{D}')$ *is a double covered graph. However, this choice for the double cover yields an irreducible double covered graph. Furthermore,* $\chi(G, \mathbf{D}') = 4 - 4 + 2 = 2$. ❑

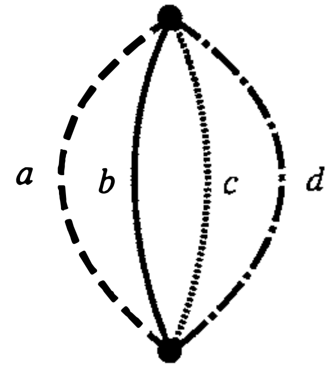

FIGURE VI.1

LEMMA VI.4. *Let* $(G, \boldsymbol{D})$ *be a double covered graph. Then* $(G, \boldsymbol{D})$ *is irreducible if and only if* $\sum \boldsymbol{D}' \neq \varnothing$ *for all proper subsets* $\boldsymbol{D}'$ *of* $\boldsymbol{D}$.

Proof. Suppose $\sum \boldsymbol{D}' = \varnothing$ for some non-empty proper subset $\boldsymbol{D}'$ of $\boldsymbol{D}$. Let $J = G[\bigcup \boldsymbol{D}']$. Hence $\boldsymbol{D}_J = \boldsymbol{D}'$. Since each edge in EJ appears twice in the sum $\sum \boldsymbol{D}'$, then $\boldsymbol{D}'$ is a double cover for J. Therefore $(J, \boldsymbol{D}')$ is a double covered graph and $(G, \boldsymbol{D})$ is reducible.

Now suppose that $(G, \boldsymbol{D})$ is reducible and let $(J, \boldsymbol{D}_J)$ be a irreducible component of $(G, \boldsymbol{D})$. Hence $\boldsymbol{D}_J \subset \boldsymbol{D}$ and by Lemma VI.2, $\sum \boldsymbol{D}_J = \varnothing$, as required. ❑

3. THE DUAL OF A DOUBLE COVERED GRAPH

Let $(G, \boldsymbol{D})$ be a double covered graph. The *dual* $G^\dagger$ *(with respect to* $\boldsymbol{D}$*)* is defined for G as follows. Let $G^\dagger$ be a graph whose vertex set is $\boldsymbol{D}$ and whose edge set is EG. Any edge e in $EG^\dagger$ is to join the two circuits in $VG^\dagger$ containing e.

LEMMA VI.5. *If* $(G, \boldsymbol{D})$ *is a double covered graph, then* $(G, \boldsymbol{D})$ *is irreducible if and only if* $G^\dagger$ *is connected.*

Proof. Suppose that $(G, \boldsymbol{D})$ is irreducible and that $G^\dagger$ is not connected. Let K be a component of $G^\dagger$. Therefore $VK \subset VG^\dagger$. Let $\boldsymbol{D}' = VK$. Evidently $\varnothing \neq \boldsymbol{D}' \subset \boldsymbol{D}$. Now, each edge in EK joins exactly two vertices, and so each edge in EG appears twice in the sum $\sum \boldsymbol{D}'$.

Hence $\sum \boldsymbol{D}' = \varnothing$, which contradicts Lemma VI.4. We conclude that $G^\dagger$ has just one component.

Now suppose that $G^\dagger$ is connected and that $(G, \boldsymbol{D})$ is reducible. By Lemma VI.4, there exists a non-empty set $\boldsymbol{D}'$ of vertices in $VG^\dagger$ such that $\sum \boldsymbol{D}' = \varnothing$ and $\boldsymbol{D}' \subset \boldsymbol{D} = VG^\dagger$. Therefore $\partial \boldsymbol{D}' = \varnothing$, which implies that $G^\dagger$ is disconnected, a contradiction. ❑

4. UNIFORM DOUBLE COVERED GRAPHS

Let $(G, \boldsymbol{D})$ be a double covered graph. Let v be a vertex in VG. Since each circuit in $\boldsymbol{D}$ which meets ∂v contains exactly two edges of ∂v, then ∂v is the union of a set $\boldsymbol{D}^\dagger(v)$ of disjoint circuits in $G^\dagger$. If $|\boldsymbol{D}^\dagger(v)| = 1$ for all $v \in VG$, then we say that $(G, \boldsymbol{D})$ is a *uniform* double covered graph. A vertex v in two irreducible components of G would violate the property $|\boldsymbol{D}^\dagger(v)| = 1$. Hence all connected uniform double covered graphs are irreducible.

5. INDEPENDENT SETS OF CYCLES

A set E of edges in a graph G *induces* a cycle C if $C \subseteq E$. Similarly, a set S of cycles in G *induces* a cycle C if $C \subseteq \bigcup S$. The members of a set S of cycles are *dependent* if there is $C \in S$ induced by $S - \{C\}$. The cycles in S are *independent* if they are not dependent.

Let $S = \{S_1, S_2, ..., S_m\}$ be an independent set of cycles. For each i, where $1 \le i \le m$, there exists an edge $e_i \in S_i$ such that

$e_i \notin \bigcup(S - \{S_i\})$. We say that e_i *represents* S_i and that $R = \{e_1, e_2, ..., e_m\}$ is a *representation* of $\boldsymbol{S}$.

6. SEPARATING CYCLES

Let $(G, \boldsymbol{D})$ be an irreducible double covered graph and let $G^\dagger$ denote the dual of G with respect to $\boldsymbol{D}$. A cycle C in G is said to *separate* $(G, \boldsymbol{D})$ if $c(G^\dagger - C) > c(G^\dagger)$. This is the Second Jordan Curve Property discussed by Vince and Little in [31]. Our purpose here is to extend this concept from a single cycle to a set of cycles to obtain a generalisation of the following theorem of [31]. This generalisation is of the same nature as the generalisation of Stahl's work ([24, 25]) which appeared in Chapter III.

Theorem VI.6. *Let $(G, \boldsymbol{D})$ be an irreducible double covered graph. Then every cycle in G separates if and only if $\chi(G, \boldsymbol{D}) = 2$.* ❑

A set S of cycles in G is said to *separate* a double covered graph $(G, \boldsymbol{D})$ if S induces a cycle that separates $(G, \boldsymbol{D})$. A set S of m independent circuits that does not separate is called an *m-fundamental set.* A set of independent semicycles is *fundamental* if it is m-fundamental for some m. The *first betti number* $h(G, \boldsymbol{D})$ of the double covered graph $(G, \boldsymbol{D})$ is the maximum size of a fundamental set in G. The next three sections are devoted to a proof of the following theorem.

THEOREM VI.7. *The maximum size of a fundamental set in an irreducible double covered graph* $(G, \mathbf{D})$ *is* $2 - \chi(G, \mathbf{D})$. ❑

7. IMPLIED CYCLES

Let G' be the graph obtained from a graph G by contracting a link $e \in EG$ to a vertex v'. Let e join v and w in G. Let C' be a circuit in G'. If $v' \notin VC'$ then evidently C' is a circuit in G, and we let $C = C'$. If $v' \in VC'$ then let e_1 and e_2 be the edges of C' incident on v'. (If C' is a loop then $e_1 = e_2$.) Without loss of generality assume that e_1 is incident on v in G. If e_2 is incident on w in G then we let $C = C' \cup \{e\}$; otherwise we let $C = C'$. In each case, we say that C is the circuit in G *implied* by C'.

Conversely if D is a cycle in G then clearly $D' = D - \{e\}$ is a cycle in G'. We say that D' is the cycle *implied* by D.

8. LINK CONTRACTION SEQUENCES

If e_1 and e_2 are two distinct links in a graph G then clearly e_2 might not be a link in the graph obtained from G by contracting e_1. Therefore we speak of contracting a *sequence* of links. Since our graphs have a finite number of edges, the link contraction sequence will always be finite.

Let $(G, \mathbf{D})$ be an irreducible double covered graph. Contract links in G one at a time until none is left. Then we have a finite sequence of graphs $G = G_n, G_{n-1}, \ldots, G_1$ where G_i is obtained by contracting a link e_i in G_{i+1} for $1 \le i \le n-1$. Clearly G_1 is a connected graph with $|VG_1| = 1$. Let $A = \{e_1, e_2, ..., e_n\}$.

LEMMA VI.8. *The set A is the edge set of a spanning tree in G.*

Proof. Evidently $VG[A] = VG$ since there are no links in G_1. If we contract all but one of the edges in a circuit, then that final edge of the circuit becomes a loop. From this observation we see that there are no circuits in A. Finally, we note that A must be connected since G is connected. ❑

LEMMA VI.9. *The graph $G^\dagger - A$ is connected.*

Proof. Let K be any component of $G^\dagger - A$. Consider the cycle $C = \sum VK$ in G. By construction $C \subseteq A$. If $C \neq \varnothing$, then A contains a circuit, contradicting the fact that A is a tree. Therefore $C = \varnothing$. It follows that $VK = VG^\dagger$ since $G^\dagger$ is connected. Hence K is the only component of $G^\dagger - A$. ❑

Since $G^\dagger - A$ is connected, it is spanned by a tree. Therefore, let T be a spanning tree of $G^\dagger - A$, and let $T' = EG - A - T$.

LEMMA VI.10. $|T'| = 2 - \chi(G, \boldsymbol{D})$ *and hence the size of any spanning tree in $G^\dagger - A$ is $|EG - A| - 2 + \chi(G, \boldsymbol{D})$.*

Proof. Now $|T'| = |EG - A - T| = |EG| - |A| - |T|$ since $A \cap T = \varnothing$. Therefore $|T'| = |EG| - |A| - |V(G^\dagger - A)| + 1$ since the size of any spanning tree for $G^\dagger - A$ is $|V(G^\dagger - A)| - 1$. Hence $|T'| = |EG| - |A| - |VG^\dagger| + 1 = |EG| - |A| - |D| + 1 = |EG| - |VG| - |D| + 2 = 2 - \chi(G, D)$. ❑

If C_1 is a circuit in G_1 then we apply the definition for the implied circuit inductively to obtain a circuit C_i in G_i. We say that C_i is the circuit in G_i *implied* by C_1. Clearly $C_i \subseteq C_1 \cup \{e_1, e_2, ..., e_{i-1}\}$. If D_n is a circuit in $G_n = G$ then we apply the definition for the implied cycle inductively to obtain a cycle D_i in G_i. We say that D_i is the cycle in G_i *implied* by D_n. Clearly $D_i = D_n - \{e_n, e_{n-1}, ..., e_i\}$.

For each $e_i \in T'$ where $i = 1, ..., 2 - \chi(G, D)$, let S_i be the circuit $\{e_i\}$ in G_1. Let $S' = \{S_1, S_2, ..., S_{2-\chi(G,D)}\}$. Since all the e_i's are distinct, S' is a set of independent circuits in G_1. Let S be the set of circuits in G that are implied by the circuits of S'. Since S' is a set of independent circuits then S is a set of independent circuits. The following lemma shows that S does not separate G.

LEMMA VI.11. *If D is a cycle in G induced by S, then D does not separate G.*

Proof. Suppose D separates G. Therefore $G^\dagger - D$ has at least two components. Let K be a component of $G^\dagger - D$, and let $L = \partial VK$. By construction, $L \subseteq D$. Also $L \cap T \neq \varnothing$ for otherwise T would not span $G^\dagger$. Therefore $T \cap D \neq \varnothing$. Since D is induced by S, $D \subseteq \bigcup S$. It follows that $T \cap (\bigcup S) \neq \varnothing$. However $T \cap T' = \varnothing$ and $T \cap A = \varnothing$.

Therefore $T \cap (\bigcup S) = T \cap (\bigcup S' \cup A) = T \cap (T' \cup A) = \varnothing$, a contradiction. ❑

From these results we make the following observation.

THEOREM VI.12. *If $(G, \boldsymbol{D})$ is an irreducible double covered graph, then $h(G, \boldsymbol{D}) \geq 2 - \chi(G, \boldsymbol{D})$.* ❑

9. AN UPPER BOUND ON $h(G, \boldsymbol{D})$

Suppose $(G, \boldsymbol{D})$ is an irreducible double covered graph. The purpose of this section is to show that any set S of independent circuits of cardinality greater than $2 - \chi(G, \boldsymbol{D})$ will separate $(G, \boldsymbol{D})$. If G is a graph with just one edge e then clearly $|VG| = 1$ and $\boldsymbol{D} = (\{e\}, \{e\})$. Furthermore, the only set of independent circuits is $\{\{e\}\}$ which separates G. Henceforth we assume that G contains at least two edges.

Let $S = \{S_1, S_2, ..., S_m\}$ be a set of m independent circuits in G, where $m > 2 - \chi(G, \boldsymbol{D})$, and let R be a representation of S. Hence $|R| = m > 2 - \chi(G, \boldsymbol{D})$. In G, contract a link e_{n-1} in $E = \bigcup S - R$. Let G_{n-1} denote the resulting graph, and let $E_{n-1} = \bigcup S - R - \{e_{n-1}\}$. Repeat this operation inductively until we obtain a graph G_j and edge set E_j such that E_j contains just loops in G_j.

LEMMA VI.13. *If e is an edge of $\bigcup S$ that is in G_j then e is a loop in G_j.*

Proof. We have that $E_j = \bigcup S - R - \{e_{n-1}, e_{n-2}, ..., e_j\}$ is a set of loops in G_j and that $EG_j = EG - \{e_{n-1}, e_{n-2}, ..., e_j\}$. Therefore it is sufficient to show that any edge $e \in R$ is a loop in G_j. However, this follows from the fact that if we contract all the edges of a circuit but one, then the final edge of that circuit becomes a loop. ❑

If $|VG_j| = 1$ then we let $G_1 = G_j$. Otherwise we contract all the remaining links from G_j until none is left and let G_1 denote the resulting graph. We have a finite sequence of graphs $G = G_n, G_{n-1}, ..., G_j, ..., G_1$ where G_i is obtained by contracting a link e_i in G_{i+1} for $1 \le i \le n-1$. Let $A = \{e_1, ..., e_{n-1}\}$. From the proof of Lemma VI.13, all the edges in R are edges in G_1. Let $T = EG - A - R$. Therefore $|T| = |EG - A| - |R| < |EG - A| - 2 + \chi(G, \boldsymbol{D})$. By Lemma VI.10, T has too few edges to be the edge set of a spanning tree for $G^\dagger - A$. Therefore $c(G^\dagger - A - R) > 1$. Let K be a component of $G^\dagger - A - R$ (possibly a single vertex). Hence $VK \subset V(G^\dagger - A - R)$. Consider the cycle $C = \sum VK$ in G. Since $(G, \boldsymbol{D})$ is irreducible, we see that if $C = \varnothing$ then $VK = VG^\dagger = V(G^\dagger - A - R)$, a contradiction. Hence by construction C is a cycle that separates $(G, \boldsymbol{D})$, and whose edges are contained in $A \cup R$. It remains to show that C is induced by $\boldsymbol{S}$.

LEMMA VI.14. *The cycle C is induced by $\boldsymbol{S}$.*

Proof. Suppose C is not induced by $\boldsymbol{S}$, and let e be an edge in C that is not in $\bigcup \boldsymbol{S}$. Since $(G, \boldsymbol{D})$ is irreducible and $|EG| > 2$, then G cannot contain any loops. Therefore e is a link in G. Let $C_1 = C - A$.

Therefore C_1 is the cycle in G_1 implied by C. Since $C \subseteq A \cup R$ then $C_1 \subseteq R$. Also $e \notin R$, for otherwise e would be in $\bigcup S$. Hence $e \in A$ and $e = e_k$ for some k. Evidently $1 \leq k < j$; otherwise $e \in \bigcup S$. We choose e so as to minimise k. Let C_{k+1} be the cycle implied by C in G_{k+1}. Hence $e \in C_{k+1} \subseteq R \cup \{e_1, e_2, ..., e_k\}$ and $\{e_1, e_2, ..., e_{k-1}\} \subseteq \bigcup S$. By Lemma VI.13 and the fact that $1 \leq k < j$, e is the only link in C_{k+1}. However, this implies that both vertices in ∂e have odd degree in C_{k+1} and hence C_{k+1} is not a cycle, a contradiction. ❑

Since C is a cycle induced by S that separates $(G, \boldsymbol{D})$, we have the following theorem.

THEOREM VI.15. *If $(G, \boldsymbol{D})$ is an irreducible double covered graph, then $h\,(G, \boldsymbol{D}) \leq 2 - \chi(G, \boldsymbol{D})$.* ❑

Theorem VI.7 now follows from Theorem VI.12 and Theorem VI.15.

10. GEM ENCODED DOUBLE COVERED GRAPHS

Let $(G, \boldsymbol{D})$ be a doubled covered graph where $\boldsymbol{D} = (D_1, D_2, ..., D_n)$. Furthermore, suppose G to have no loops. We form a 3-graph K in the following way. The vertices of K are the ordered triples of the form (v, e, i), where e is an edge of D_i incident on vertex v. Two vertices of K are adjacent if and only if they differ in exactly one component. The unique edge joining them is coloured blue, yellow or red according to whether that component is the first, second or third. If $(v, e, i) \in VK$, then exactly one edge of $D_i - \{e\}$ is incident on v, e is incident on exactly one vertex of $VG - \{v\}$, and $e \in D_j$ for exactly one integer $j \neq i$. We deduce that K is a 3-graph. Moreover the red-blue bigons are squares and so K is a gem. This gem is said to *encode* $(G, \boldsymbol{D})$. The set of all edges of K joining vertices with a given jth component, where $j \in \{1, 2, 3\}$, constitutes a single bigon or a union of bigons of the same type.

LEMMA VI.16. *The set Y of all edges in EK joining vertices with a given second component e constitutes a single red-blue bisquare.*

Proof. Let D_i and D_j denote the two circuits in $\boldsymbol{D}$ that contain e. Let e join v and w. Then $VY = \{(v, e, i), (v, e, j), (w, e, j), (w, e, i)\}$. Hence Y is a single red-blue bisquare. ❑

LEMMA VI.17. *The set R of all edges in EK joining vertices with a given third component D constitutes a single blue-yellow bigon.*

Proof. Let $D = \{e_1, e_2, ..., e_n\}$ where e_j is adjacent to e_{j+1} whenever $1 \leq j < n$. (It follows that e_n is adjacent to e_1.) Let v_j denote the vertex in VG incident on both e_j and e_{j+1}. Let v_n denote the vertex incident on both e_n and e_1.

Now consider the vertex set $X = \{(v_1, e_1, i), (v_2, e_1, i), (v_2, e_2, i), (v_3, e_2, i), ..., (v_n, e_n, i), (v_1, e_n, i)\}$ in K. Since the vertices in X differ in only their first and second components, X is the vertex set of a single blue-yellow bigon. Moreover, since v_j is incident on two edges in D_i, then $X = VR$. Hence R is a single blue-yellow bigon. ❑

LEMMA VI.18. *The set B of all edges in K joining vertices with a given first component v is a single red-yellow bigon for all $v \in VG$ if and only if $(G, \boldsymbol{D})$ is a uniform double covered graph.*

Proof. Clearly the number of circuits in $\boldsymbol{D}$ that pass through v is $\deg_G(v)$. Let D_0 be such a circuit. Let e_0 and e_1 denote the two edges of D_0 incident on v. Let D_1 be the circuit in $\boldsymbol{D}$ such that $D_1 \neq D_0$ and $e_1 \in D_1$. Let e_2 denote the edge of D_1 other than e_1 incident on v. Proceeding inductively, we obtain a set $\boldsymbol{D}' = \{D_0, D_1, \ldots, D_{m-1}\}$ of circuits and a set $L = \{e_0, e_1, \ldots, e_{m-1}\}$ of edges such that $D_i \cap \partial v = \{e_i, e_{i+1}\}$ and $D_i \cap D_{i+1} \cap \partial v = \{e_i\}$ for all i, where the subscripts are read modulo m. Now consider the vertex set $X = \{(v, e_0, 0), (v, e_1, 0), (v, e_1, 1), (v, e_2, 1), \ldots, (v, e_{m-1}, m-1), (v, e_{m-1}, 0)\}$. Since the vertices in X differ in only their second and third components, X is the vertex set of a red-yellow bigon B'. If $m < \deg_G(v)$, then L would be a circuit in $G^\dagger$ properly included in $\boldsymbol{D}^\dagger(v)$. This would imply that $|\boldsymbol{D}^\dagger(v)| > 1$. Hence we conclude that $m = \deg_G(v)$ for all vertices $v \in VG$ if and only if $(G, \boldsymbol{D})$ is a uniform double covered graph. In the case where $m = \deg_G(v)$, we see that $L = \partial v$, and that $\boldsymbol{D}'$ is the set of all circuits in $\boldsymbol{D}$ passing through v.

From this observation, we see that X is the set of all vertices of K with first component v. Hence $B = B'$, and B is a single red-yellow bigon, as required. ❑

LEMMA VI.19. *Let* $(G, \boldsymbol{D})$ *be a uniform double covered graph. Let* K *be the gem that encodes* $(G, \boldsymbol{D})$*. Then* $\chi(G, \boldsymbol{D}) = \chi(K)$*.*

Proof. By Lemmas VI.16, VI.17 and VI.18, $|\boldsymbol{Y}(K)| = |EG|$, $|\boldsymbol{R}(K)| = |\boldsymbol{D}|$ and $|\boldsymbol{B}(K)| = |VG|$. It follows that the number of bigons in K is $|EG| + |VG| + |\boldsymbol{D}|$. Hence

$$\begin{aligned} \chi(K) &= r(K) - \frac{|VK|}{2} \\ &= |EG| + |VG| + |\boldsymbol{D}| - 2|EG| \\ &= \chi(G, \boldsymbol{D}). \ ❑ \end{aligned}$$

11. REFINED 3-GRAPHS

Let K be a 3-graph. We say K is *red-refined* if all blue-yellow bigons are semicycles. Similarly, K is *yellow-refined* if all red-blue bigons are semicycles. If K is both red-refined and yellow-refined, then K is said to be refined. In Chapter VII, we show how one may generate a refined 3-graph from an arbitrary 3-graph by dipole cancellations and creations.

LEMMA VI.20. *Let* K *be the gem that encodes a uniform double covered graph* $(G, \boldsymbol{D})$*. If* R *is a blue-yellow bigon in* K*, then* R *is a semicycle. Hence* K *is a red-refined graph.*

Proof. Let i be the common third component of the vertices in VR. Suppose R is not a semicycle. Therefore R must meet a red-yellow bigon, B say, in more than one yellow edge. Let v be the common first component of the vertices in VB. Let c and c' denote two distinct yellow edges in $R \cap B$. Let (v, e_1, i) and (v, e_2, i) be the two vertices incident on c. Similarly let (v, e_3, i) and (v, e_4, i) be the two vertices incident on c'. Since c and c' are distinct, e_1, e_2, e_3 and e_4 are all distinct. However, this implies that e_1, e_2, e_3 and e_4 are four edges in C incident on v in G. This contradicts the fact that C is a circuit. The lemma follows. ❑

LEMMA VI.21. *Let K be the gem that encodes a uniform double covered graph (G, D) with no loops. Then K is a yellow-refined 3-graph.*

Proof. Let Y be a red-blue bigon in K. By the argument in Lemma VI.16, the four vertices in VY do not have a common first component. Thus Y meets a given red-yellow bigon in at most one edge. Hence K is a yellow-refined graph. ❑

By Lemma VI.18, G clearly underlies K if (G, D) is a uniform double covered graph. Hence a circuit in G corresponds to a semicycle in K and a semicycle in K represents a circuit in G or a single edge in G.

12. RED-YELLOW REDUCTIONS

Let G be the graph obtained from a 3-graph K by contracting the red-yellow bigons to single vertices. We say that G is the *red-yellow reduction* of K. Clearly $c(G) = c(K)$. If S is a semicycle in K, then the requirements that S should be a circuit that has a blue edge, and that each bead should have just two poles, reveal that $\beta(S)$ is a circuit in G. We say that $\beta(S)$ *corresponds* to S. On the other hand, if C is a circuit of G, then C is the blue edge set of a family of semicycles in K. We say that each semicycle in this family *represents* C. If K and L are congruent 3-graphs then evidently the red-yellow reductions of K and L are the same graph.

EXAMPLE VI.22. *If K is the* 3-*graph of Figure* III.3 *then the graph G shown in Figure* VI.2 *is the red-yellow reduction of K.* ❑

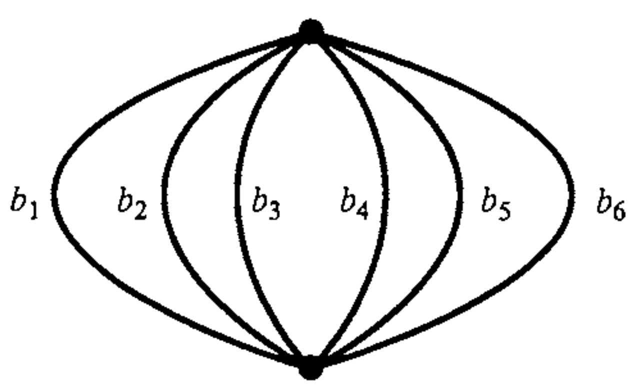

FIGURE VI.2

13. OBTAINING DOUBLE COVERED GRAPHS FROM 3-GRAPHS

Let K be a refined 3-graph, and let G be the red-yellow reduction of K. Label the red-yellow bigons of K with $Y_1, Y_2, ..., Y_n$ and the blue-yellow bigons of K with $R_1, R_2, ..., R_m$. Clearly $D = (\beta(Y_1), \beta(Y_2), ..., \beta(Y_n), \beta(R_1), ..., \beta(R_n))$ is a circuit double cover of K. We say that (G, D) is the *double covered reduction* of K.

LEMMA VI.23. *Let (G, D) be the double covered reduction of a refined 3-graph K. Then (G, D) is a uniform double covered graph.*

Proof. Let v be a vertex in VG. Therefore v corresponds to a unique red-yellow bigon B in K such that $\partial v = \partial B$. Let $U_0, U_1, ..., U_{n-1}$ denote the bigons in K that meet B. For the present proof, we assume all subscripts are read modulo n. Since K is refined, $|U_i \cap B| = 1$ for all i. Hence we may assume that these bigons are labelled so that $(U_i \cup U_{i+1}) \cap B$ is a path of length 2 in B. Let b_i be the blue edge adjacent to both edges in $(U_i \cup U_{i+1}) \cap B$, for all i. Hence $\partial v = \partial B = \{b_0, b_1, .., b_{n-1}\}$. For each i, $\beta(U_i)$ is a vertex in $G^\dagger$. Since $U_{i-1} \cap U_i \cap \partial B = \{b_{i-1}\}$ and $U_i \cap U_{i+1} \cap \partial B = \{b_i\}$ for all i, we conclude that ∂v is a circuit in $G^\dagger$ with vertex set $\{\beta(U_0), ..., \beta(U_{n-1})\}$. Hence $|D^\dagger(v)| = 1$, as required. ❑

LEMMA VI.24. *Let (G, D) be the double covered reduction of a refined 3-graph K. Then $\chi(G, D) = \chi(K)$.*

Proof. Clearly $|D| = |Y(K)| + |R(K)|$, $|EG| = |\beta(K)| = \frac{|VK|}{2}$ and $|VG| = |B(K)|$. Therefore

$$\begin{aligned}\chi(G, D) &= |Y(K)| + |R(K)| - \frac{|VK|}{2} + |B(K)| \\ &= \chi(K). \square\end{aligned}$$

14. RELATING THE SEPARATION PROPERTIES

Theorem VI.25 below states that it is possible to translate back and forth between the language of uniform double covered graphs and cycles and the language of 3-graphs and b-cycles in such a way that special cases of Theorems III.15 and VI.7 are merely restatements of the same result in different terminology.

THEOREM VI.25. *The maximum size of a fundamental set in a connected refined* 3*-graph K is* $2 - \chi(K)$ *if and only if the maximum size of a fundamental set in a connected uniform double covered graph* (G, D) *is* $2 - \chi(G, D)$.

The proof of this theorem is given in Lemmas VI.26 and VI.27.

LEMMA VI.26. *If the maximum size of a fundamental set in a connected uniform double covered graph* (G, D) *is* $2 - \chi(G, D)$ *then the maximum size of a fundamental set in a connected refined* 3*-graph K is* $2 - \chi(K)$.

Proof. Let K be a connected refined 3-graph and suppose $h(G, D) = 2 - \chi(G, D)$ for all connected uniform double covered graphs. Let (G, D) be a double covered reduction of K. Hence $\chi(G, D) = \chi(K)$.

Firstly, let $S = \{S_1, S_2, ..., S_{2-\chi(G, D)}\}$ be a fundamental set in (G, D). Let $S' = \{S_1', S_2', ..., S_{2-\chi(K)}'\}$ where S_i' is a semicycle that represents S_i for each i. Hence $\beta(S_i') = S_i$ for all i. Clearly S' is b-independent since S is independent. We shall show that S' is a fundamental set. Suppose that S' induces a separating b-cycle. Then it induces a separating cycle of the form $\sum U$ for some set $U \subseteq Y(K) \cup R(K)$. Let V be the set of circuits in D that correspond to the circuits in U. Clearly $\sum V = \beta(\sum U) \subseteq \beta(\bigcup S') \subseteq \bigcup S$. Therefore $\sum V$ is a separating cycle induced by S, which contradicts the fact that S is fundamental. Hence we conclude that S' is a fundamental set in K. Thus the maximum size of a fundamental set in K is no less than $2 - \chi(K)$.

Now, let $S' = \{S_1', S_2', ..., S_m'\}$ be a b-independent set of semicycles in K, where $m > 2 - \chi(K)$. Then $S = \{\beta(S_1'), \beta(S_2'), ..., \beta(S_m')\}$ is an independent set of circuits in G. Since $m > \chi(G, D)$, S separates (G, D). Hence S induces a cycle of the form $\sum V$ for some set $V \subseteq D$. Let U be the set of bigons in K that correspond to the circuits in V. Clearly $\beta(\sum U) = \sum V \subseteq \bigcup S = \beta(\bigcup S')$, and therefore $\sum U$ is a separating cycle induced by S'. Hence S' separates, and we conclude that the maximum size of a fundamental set in K is no more than $2 - \chi(K)$, as required. ❑

LEMMA VI.27. *If the maximum size of a fundamental set in a connected refined* 3*-graph* K *is* $2-\chi(K)$ *then the maximum size of a fundamental set in a connected uniform double covered graph* $(G, \boldsymbol{D})$ *is* $2-\chi(G, \boldsymbol{D})$.

Proof. Let $(G, \boldsymbol{D})$ be a connected uniform double covered graph and suppose that the maximum size of a fundamental set in all connected refined 3-graphs K is $2-\chi(K)$. Let K be the refined gem that encodes $(G, \boldsymbol{D})$. Hence by Lemma VI.19, $\chi(G, \boldsymbol{D}) = \chi(K)$.

Firstly, let $\boldsymbol{S} = \{S_1, S_2, ..., S_{2-\chi(K)}\}$ be a fundamental set in K. Since $\boldsymbol{S}$ does not separate, no semicycle in $\boldsymbol{S}$ can correspond to an edge in G, for otherwise $\boldsymbol{S}$ would induce a red-blue bisquare. Let $\boldsymbol{S}' = \{S_1', S_2', ..., S_{2-\chi(G,\, D)}'\}$ where S_i' is the circuit in G that underlies S_i for each i. If both blue edges of a red-blue bigon in K were in $\bigcup\boldsymbol{S}$, then $\boldsymbol{S}$ would separate. Hence we conclude that $\boldsymbol{S}'$ is independent since $\boldsymbol{S}$ is b-independent. We shall show that $\boldsymbol{S}'$ is a fundamental set. Suppose that $\boldsymbol{S}'$ induces a separating cycle. Then it induces a cycle of the form $\sum\boldsymbol{V}$ for some set $\boldsymbol{V} \subseteq \boldsymbol{D}$. Let $\boldsymbol{U}$ be the set of blue-yellow bigons in K that correspond to the circuits in $\boldsymbol{V}$. Let $\boldsymbol{Y}_1$ be the set of red-blue bigons in K that meet $\bigcup\boldsymbol{U}$ in two blue edges. Let $\boldsymbol{Y}_2$ be the set of red-blue bigons in K that have one blue edge in $\bigcup\boldsymbol{U}$ and the other in $\bigcup\boldsymbol{S}$. Hence $\boldsymbol{Y}_1$ and $\boldsymbol{Y}_2$ are disjoint. Evidently $D = \sum(\boldsymbol{U} \cup \boldsymbol{Y}_1 \cup \boldsymbol{Y}_2)$ is induced by $\boldsymbol{S}$. Furthermore D is a separating cycle since it is a sum of bigons. This contradicts the fact that S does not separate. Therefore we conclude that S' is fundamental, and that $h(G, D) \geq 2 - \chi(G, D)$.

Now, let $S' = \{S_1', S_2', ..., S_m'\}$ be an independent set of circuits in G, where $m > 2 - \chi(G, D)$. Let $S = \{S_1, S_2, ..., S_m\}$, where S_i is a semicycle in K that represents S_i' for each i. Furthermore, we may assume that S is chosen so that no red-blue bisquare in K has both its blue edges in $\bigcup S$. Clearly S is a b-independent set of semicycles in K. Since $m > \chi(K)$, S separates K. Hence S induces a cycle of the form $\sum U$ for some set $U \subseteq Y(K) \cup R(K)$. Let V be the set of circuits in D that correspond to the blue-yellow bigons in U. Let e be an edge in $\sum V$. Then e corresponds to a red-blue bigon Y. We claim that exactly one blue edge of Y belongs to $\bigcup S$, and therefore $e \in \bigcup S'$. Indeed, exactly one blue edge of Y belongs to $\bigcup(U \cap R(K))$, and hence $\bigcup S$, since $Y \notin U \cap Y(K)$. It follows that $\sum V \subseteq \bigcup S'$, which implies that S' separates. Moreover, $\beta(\sum U) = \sum V \subseteq \bigcup S = \beta(\bigcup S')$, and therefore $\sum U$ is a separating cycle induced by S'. Hence S' separates, and we conclude that $h(G, D) \leq 2 - \chi(G, D)$, as required. ❑

Chapter VII

MAC LANE'S THEOREM FOR 3-GRAPHS

1. INTRODUCTION

MacLane, in an attempt to make a partial separation between graph theory and topology, endeavoured to prove that a given graph can be imbedded in the sphere if and only if it had a certain combinatorial property. However the proof required topological arguments. In this chapter we present a purely combinatorial proof of MacLane's theorem which evolves from the tools we have developed in the preceeding chapters.

Recall that a connected 3-graph K is planar if $\chi(K) = 2$. More generally, a 3-graph K is *planar* if each component is planar. Hence any 3-graph K is planar if $\chi(K) = 2c(K)$. A graph G is *planar* if it underlies a planar gem.

A family of cycles in a graph G is said to be a *spanning family* if its components span $\mathbf{Z}(G)$. Similarly we may talk about a *spanning set* of cycles of G. If the cycles of the spanning family or spanning set constitute a cycle double cover, then this double cover is also described as *spanning*. Similarly we also talk about a *spanning circuit double cover.*

THEOREM VII.21 [MAC LANE]. *A graph is planar if and only if its foundation has a spanning circuit double cover.*

2. REFINEMENTS

Suppose R to be a blue-yellow bigon of length 4 in a 3-graph K. Label the edges and vertices incident on R as in Figure VII.1a. If b is a blue 1-dipole then let $K' = K - [b]$. (See Figure VII.1b.) Let c' denote the yellow edge of K' that joins v' and w'. If $\{b', c'\}$ is a blue-yellow 2-dipole in K' then let $K'' = K' - [b', c']$. (See Figure VII.1c.) We say that K'' is obtained from K by *cancellation* of the blue-yellow bigon R. Let a and a' denote the distinct red edges that join v_1 and w_1, and v_2 and w_2, respectively. The inverse operation is described as *splitting* a and a' to *create* the blue-yellow bigon R.

a)

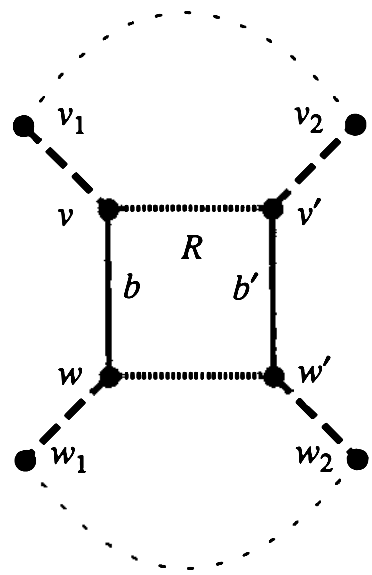

b)

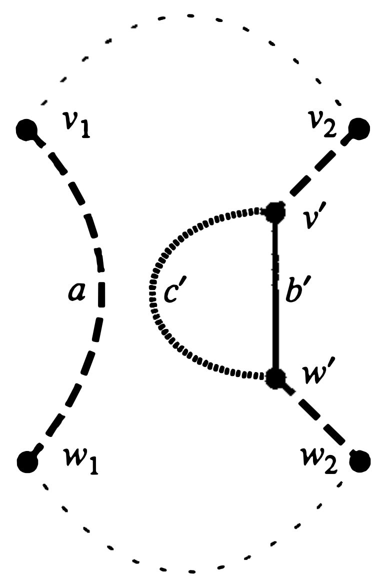

c)

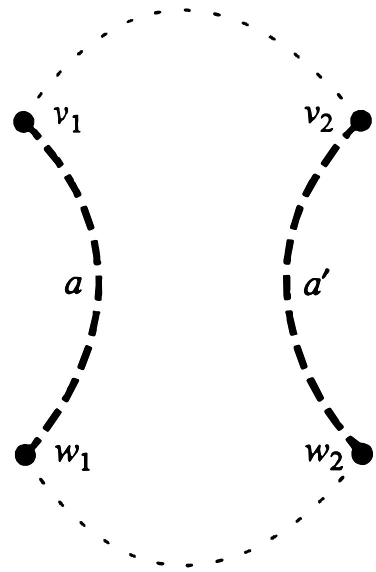

FIGURE VII.1

Let K be a 3-graph. Suppose there is a blue-yellow bigon R in K that is not a semicycle. Then there is a red-yellow bigon B such that $B \cap R$ contains two yellow edges c and c'. Hence $B - \{c, c'\}$ is the disjoint union of two paths P_1 and P_2 each of which contains a red edge. Let a be a red edge in P_1 and let a' be a red edge in P_2. Split a and a' to create the blue-yellow bigon R' and let K' denote the resulting graph. The red-yellow bigon B has now become two red-yellow bigons A and A' in K'. Furthermore R is a blue-yellow bigon in K' that meets A in fewer edges than R meets B. A similar statement holds for A'. We also note that R' is a semicycle in K', and that any blue-yellow bigons in K that are semicycles in K are semicycles in K'. Proceeding inductively we obtain a 3-graph such that R meets a red-yellow bigon in at most one edge, and hence is a semicycle. We conclude that it is possible to construct a 3-graph L from K by blue 1-dipole and blue-yellow 2-dipole creations such that all blue-yellow bigons in L are semicycles. Hence L is a red-refined 3-graph. We say that L is a *red-refinement* of K.

EXAMPLE VII.1. *Consider the gem K in Figure* VII.2*a. Figure* VII.2*b illustrates a red-refinement of K.* ❑

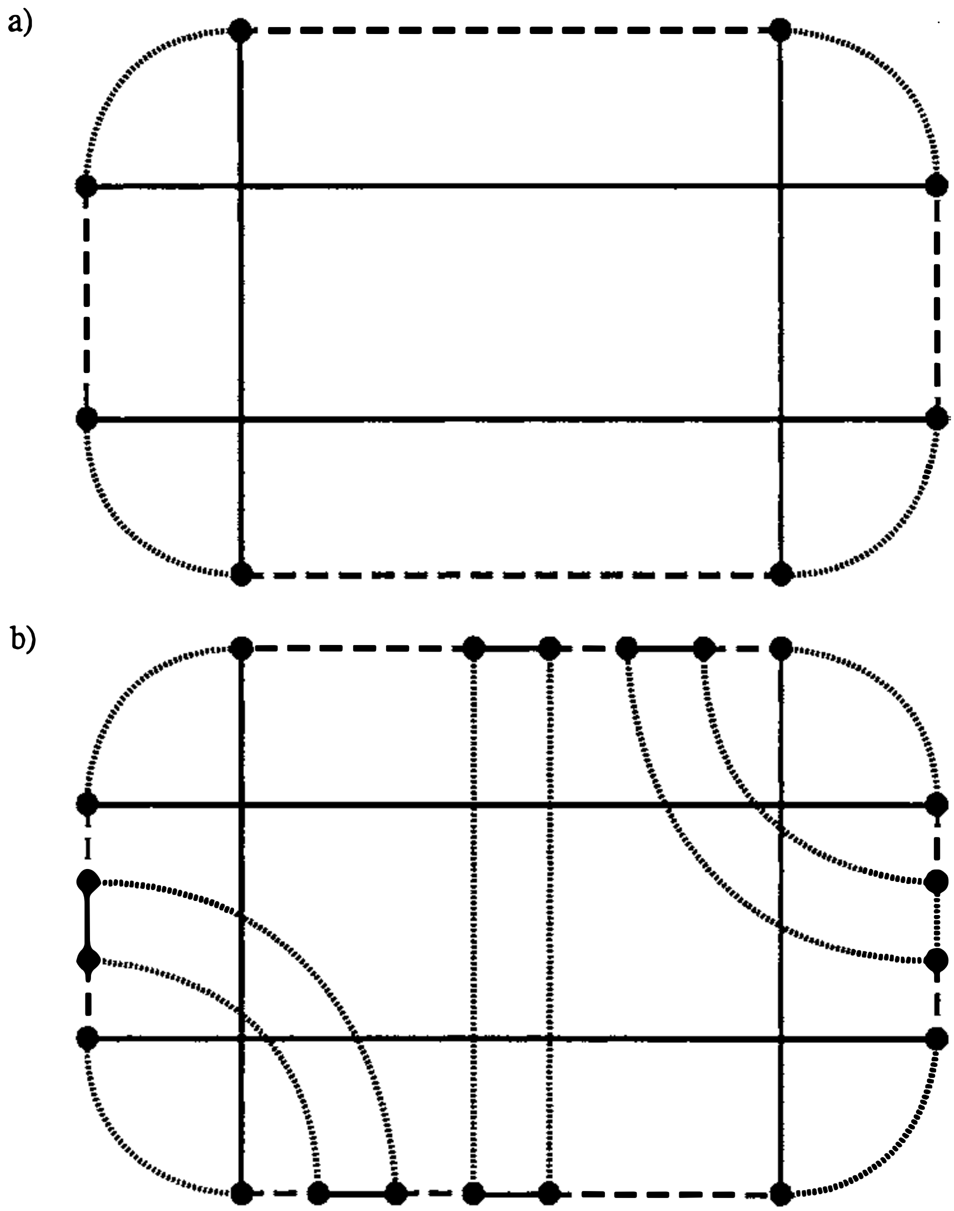

FIGURE VII.2

3. COVERS IN 3-GRAPHS

A set S of b-cycles in a 3-graph K is said to be a *b-cycle cover* if the set of intersections of members of S with $\beta(K)$ is a partition of $\beta(K)$. Similarly we talk about a *semicycle cover.* Clearly the sets $\boldsymbol{R}(K)$ and $\boldsymbol{Y}(K)$ are b-cycle covers of K. Furthermore $\boldsymbol{R}(K)$ is a semicycle cover if K is red-refined.

A set S of b-cycles in a 3-graph K is said to be a *spanning set* if $\boldsymbol{B}(K) \cup \boldsymbol{Y}(K) \cup \boldsymbol{S}$ spans the cycle space $\boldsymbol{Z}(K)$ of K. Similarly we also talk about a *spanning semicycle cover.*

EXAMPLE VII.2. *Consider the gem K of Figure* III.3. *Let* $C_1 = \{b_1, c_1, a_3, c_2, b_4, c_6\}$, $C_2 = \{b_3, a_3, c_2, a_4, b_6, c_5\}$ *and* $C_3 = \{b_5, c_2, a_4, c_3, b_2, c_4\}$. *One can easily check that* $\boldsymbol{B}(K) \cup \boldsymbol{Y}(K) \cup \{C_1, C_2, C_3\}$ *spans the cycle space of K. Therefore* $\boldsymbol{S} = \{C_1, C_2, C_3\}$ *is a spanning semicycle cover.* ❑

LEMMA VII.3. *Let C and D be two cycles in a* 3*-graph K such that $\beta(C) = \beta(D)$. Then $C = D + \bigcup \boldsymbol{B}$ for some set $\boldsymbol{B}$ of red-yellow bigons in $N(C) = N(D)$.*

Proof. $C + D$ is a cycle A which does not meet $\partial_K VB$ for any bead B of C. Hence $A \cap B \in \{\varnothing, B\}$. If $A \cap B = \varnothing$ then $C \cap B = D \cap B$, and if $A \cap B = B$ then $C \cap B = (D + B) \cap B$. Since $\beta(C) = \beta(D)$, it follows that $C = D + \bigcup \boldsymbol{B}$, where $\boldsymbol{B}$ is the set of beads B of C for which $A \cap B = B$. ❑

4. BOUNDARY COVERS

A semicycle cover S is a *boundary cover* if $B(K) \cup Y(K) \cup S$ spans the boundary space of K.

LEMMA VII.4. *Let K be a 3-graph with blue 1-dipole b and let $L = K - [b]$. There exists a boundary cover in L if there exists a boundary cover in K.*

Proof. The following uses the notation of Figure II.2. Let S denote a boundary cover of K. Let C_2 denote the semicycle in S that contains b. Suppose there exists a semicycle $C \in S$ which contains c_1. Then clearly $C_1 = C + A$ is a semicycle such that $c_1 \notin C_1$, and $c_2 \in C_1$ if and only if $c_2 \in C$. Furthermore $S_1 = (S - \{C\}) \cup \{C_1\}$ is a boundary cover of K. A similar argument may be applied to c_2. Therefore we assume that no semicycle in S contains c_1 or c_2. In particular, C_2 includes $\{a_1, b, a_2\}$.

Let $S_2 = (S - \{C_2\}) \cup \{C_3\}$ where $C_3 = (C_2 - \{a_1, b, a_2\}) \cup \{a\}$. Evidently the set of intersections of members of S_2 with $\beta(L)$ is a partition of $\beta(L)$. We claim that $B(L) \cup Y(L) \cup S_2$ spans the boundary space of L. It is sufficient to show that any blue-yellow bigon R of L is a sum of circuits in $B(L) \cup Y(L) \cup S_2$. If $c \in R$ let $R_1 = (R - \{c\}) \cup \{c_1, b, c_2\}$; otherwise let $R_1 = R$. Then R_1 is a blue-yellow bigon in K, and hence $R_1 = \sum U$ for some set U of circuits in $B(K) \cup Y(K) \cup S$. We claim that $A \in U$ if and only if $B \in U$. However this is clear, for $c_2 \in R_1$ if and only if $c_1 \in R_1$ and no semicycle in S contains either c_1 or c_2. Let Y denote the red-blue bigon in K that includes $\{a_1, b, a_2\}$ and let Y' denote the red-blue bigon in L that contains a. Let U_1 be the set obtained from U by replacing C_2 with C_3, A and B with A', and Y with Y' if necessary. Evidently $R = \sum U_1$ and $U_1 \subseteq B(L) \cup Y(L) \cup S_2$, as required.

Clearly C_3 is a semicycle in L. If S_2 is a set of semicycles, then we are done. Suppose there is a circuit D in S_2 that is not a semicycle in L. Then A' has 4 poles with respect to D, for otherwise D is not a semicycle in K. Let v, w, x, y denote the 4 poles, where $A'_y[v,w] \subseteq D$ and $x \in A'_v[w,y]$. Thus $A'_w[x,y] \subseteq D$. Clearly $C_4 = D_w[v, y] \cup A'_w[v,y]$ and $C_5 = D_v[w, x] \cup A'_v[w,x]$ are semicycles in L. Since $\beta(D)$ is the disjoint union of $\beta(C_4)$ and $\beta(C_5)$, then by Lemma VII.3, $C_4 + C_5 = D + \bigcup B$ for some set B of red-yellow bigons in L. We conclude that $S_3 = (S_2 - \{D\}) \cup \{C_4, C_5\}$ is a set of circuits, with two more semicycles than S_2, such that $B(L) \cup Y(L) \cup S_3$ spans the boundary space of L. Moreover, the set of intersections of members of S_3 with $\beta(L)$ is a partition of $\beta(L)$. Proceeding inductively we obtain a boundary cover of L. ❑

LEMMA VII.5. *Let K be a 3-graph with blue-yellow 2-dipole $\{b, c\}$ and let $L = K - [b, c]$. Then there exists a boundary cover in L if there exists a boundary cover in K.*

Proof. The following uses the notation of Figure VII.3. Let A be the red-yellow bigon in K that includes $\{a_1, c, a_2\}$. Let S denote a boundary cover of K. Let C denote the semicycle in S that contains b. Suppose $C \neq \{b, c\}$. Then clearly $\{b, c\} = C + A$. Furthermore $S_1 = (S - \{C\}) \cup \{C + A\}$ is a boundary cover of K. Therefore we assume that the semicycle C in S that contains b is $\{b, c\}$.

Let $S' = S - \{C\}$. Suppose there exists a semicycle $C_1 \in S'$ which contains c. Then clearly $C_2 = C_1 + A$ is a semicycle such that $c \notin C_2$. Furthermore $S_1 = (S - \{C_1\}) \cup \{C_2\}$ is a boundary cover of K. Therefore we assume that no semicycle in S' contains c. Hence S' is a set of semicycles in L.

Evidently the set of intersections of members of S' with $\beta(L)$ is a partition of $\beta(L)$. We claim that $B(L) \cup Y(L) \cup S'$ spans the boundary space of L, and hence S' is a boundary cover of L. It is sufficient to show that any blue-yellow bigon R of L is a sum of circuits in $B(L) \cup Y(L) \cup S'$. Clearly R is a blue-yellow bigon in K, and hence $R = \sum U$ for some set U of circuits in $B(K) \cup Y(K) \cup S$. Let Y denote the red-blue bigon in K that contains b. Since no semicycle in S' contains b or c, then evidently $b \notin \bigcup(U - \{C, Y\})$ and $c \notin \bigcup(U - \{A, C\})$. Therefore, the fact that $R \neq \{b, c\}$ implies that either $\{A, C, Y\} \subseteq U$ or $\{A, C, Y\} \cap U = \varnothing$.

Let A' and Y' denote the sets of red-yellow and red-blue bigons respectively in L that contain a. If $\{A, C, Y\} \subseteq U$, then let $U' = (U - \{A, C, Y\}) \cup \{A', Y'\}$; otherwise let $U' = U$. Evidently $R = \sum U'$ and $U' \subseteq B(L) \cup Y(L) \cup S'$, as required. ❑

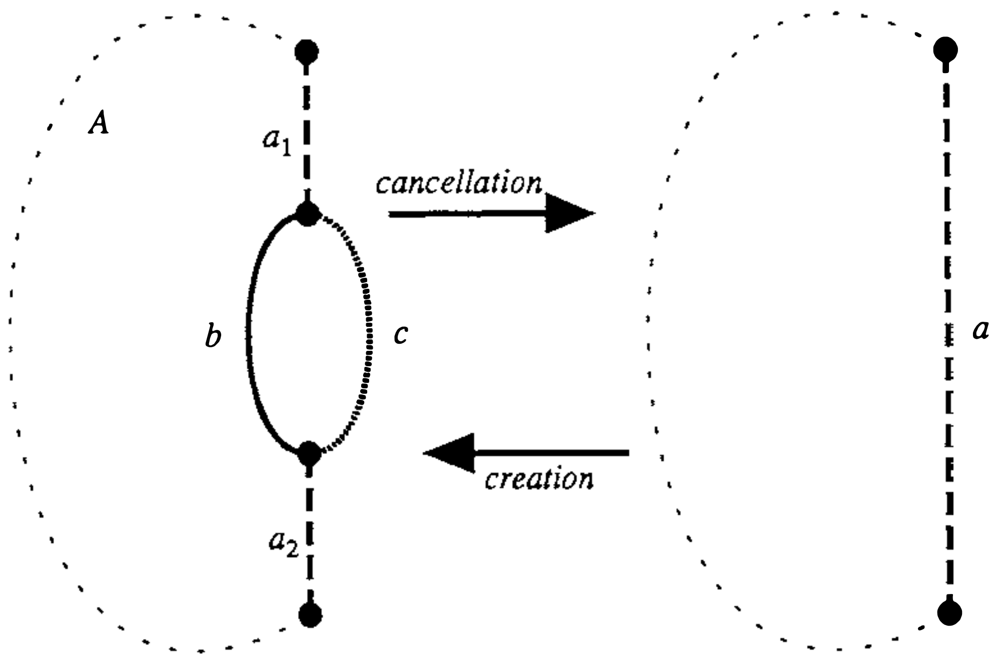

FIGURE VII.3

THEOREM VII.6. *Every* 3-*graph has a boundary cover.*

Proof. Let L be a red-refinement of an arbitrary 3-graph K. Therefore $\boldsymbol{R}(L)$ is a set of semicycles in L. It is immediate that $\boldsymbol{R}(L)$ is a boundary cover of L since the boundary space is the space spanned by $\boldsymbol{B}(L) \cup \boldsymbol{Y}(L) \cup \boldsymbol{R}(L)$. Hence our theorem follows from Lemmas VII.4 and VII.5, since K is obtained from L by a finite sequence of blue 1-dipole and blue-yellow 2-dipole cancellations. ❑

5. PARTIAL CONGRUENCE AND FAITHFULNESS

We now generalise the concept of congruence that was introduced in Chapter III. Let K and L be two 3-graphs. Suppose there exist a partition Q of $\boldsymbol{B}(L)$ and bijections θ, φ, σ between Q

and $\boldsymbol{B}(K)$, $\beta(L)$ and $\beta(K)$ and $\rho(L)$ and $\rho(K)$ respectively. Furthermore, suppose that

i) for any cell B of Q and any red edge $a \in \bigcup B$ we have $\sigma(a) \in (B)$, and

ii) for any blue edge b adjacent to a red edge a we have $\varphi(b)$ adjacent to $\sigma(a)$.

Then we say L is *partially congruent* to K (with respect to the partition Q.)

If K and L are partially congruent 3-graphs and $E = \{e_1, e_2, ..., e_n\}$ is a set of red (blue) edges in K, then for sake of conciseness we usually write E for $\sigma(E)$ ($\varphi(E)$) and e_i for $\sigma(e_i)$ ($\varphi(e_i)$) when no ambiguity results.

Let L be partially congruent to a 3-graph K and let S be a set of cycles in K. If for each $C \in S$ there exists a cycle D in L such that $\beta(D) = \beta(C)$, then we say that L is *faithful* to K (with respect to S.) The cycles D and C are said to *correspond* to each other.

EXAMPLE VII.7. *Consider the* 3*-graphs* K *and* L *in Figure* VII.8. *Let* $B_1 = \{a_1, c_3\}$, $B_2 = \{a_4, c_4\}$ *and* $B_3 = \{a_2, a_3, c_1, c_2\}$. *Let* Q *be the partition* $\{\{B_1, B_2\}, \{B_3\}\}$ *of* $\boldsymbol{B}(L)$. *From the labelling of the edges of* K, *it is clear that* L *is partially congruent to* K *with respect to the partition* Q. *Let* C *be the cycle* $\{b_2, c_4, b_3, c_1\}$ *in* K. *Since* $|\partial B_1 \cap \beta(C)| = 1$, *then* $\beta(C)$ *is not the blue edge set of a cycle in* L. *Hence* L *is not faithful to* K *with respect to* $\{C\}$. ❑

L

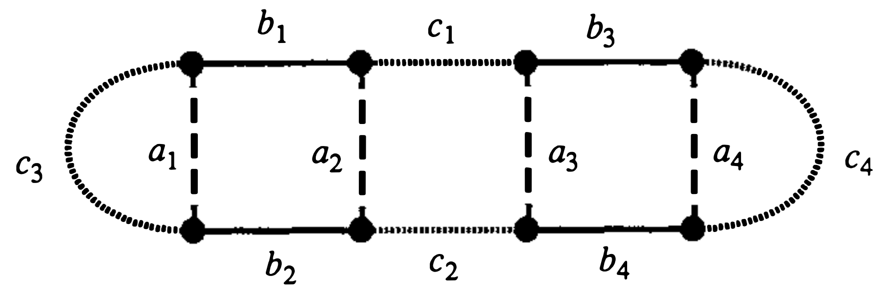

K

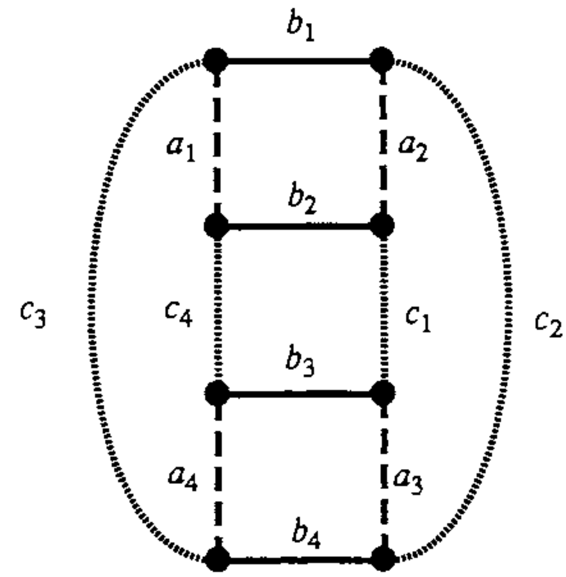

FIGURE VII.8

LEMMA VII.9. *Let L be partially congruent to a* 3*-graph K. Let C be a cycle in L. Then* $\beta(C)$ *is the blue edge set of a cycle in K.*

Proof. This follows from the fact that $|\beta(C) \cap \partial VB|$ is even for all $B \in B(L)$, and therefore $|\beta(C) \cap \partial VB|$ is even for all $B \in B(K)$. ❑

LEMMA VII.10. *Let L be a* 3*-graph partially congruent to a* 3*-graph K with respect to a partition Q. If L is faithful to K with* *respect to a spanning b-cycle cover* $\boldsymbol{S}$, *then no two red-yellow bigons in a cell of* $\boldsymbol{Q}$ *belong to the same component of* L.

Proof. We shall prove the contrapositive. Suppose B_1 and B_2 are red-yellow bigons in a cell $\boldsymbol{B}$ of $\boldsymbol{Q}$ that belong to the same component L_1 of L. Let B denote the red-yellow bigon in K that corresponds to $\boldsymbol{B}$. Since L_1 is connected, then there exists a semipath P with one terminal vertex in VB_1 and the other terminal vertex in VB_2. Hence we have $|P \cap \partial VB_1| = |P \cap \partial VB_2| = 1$, and $|P \cap \partial VB_i|$ even for all $B_i \in \boldsymbol{B}(L) - \{B_1, B_2\}$. Therefore $|\beta(P) \cap \partial VB'|$ is even for all red-yellow bigons B' in K, since B_1 and B_2 both correspond to a single red-yellow bigon. Hence $\beta(P)$ is the blue edge set of a cycle in K.

Since $\boldsymbol{S}$ is a spanning b-cycle cover, then $\beta(P)$ is the blue edge set of a cycle $\Sigma(\boldsymbol{Y} \cup \boldsymbol{U})$ for sets $\boldsymbol{Y} \subseteq \boldsymbol{Y}(K)$ and $\boldsymbol{U} \subseteq \boldsymbol{S}$. Let $\boldsymbol{R}$ be a set of cycles in L that correspond to the cycles in $\boldsymbol{U}$. Evidently, $\beta(P) = \beta(\Sigma(\boldsymbol{Y} \cup \boldsymbol{R}))$, and $\Sigma(\boldsymbol{Y} \cup \boldsymbol{R})$ is a cycle in L. Hence $|P \cap \partial VB|$ is even for all $B \in \boldsymbol{B}(L)$, a contradiction to the fact that $|P \cap \partial VB_1| = |P \cap \partial VB_2| = 1$. ❑

EXAMPLE VII.11. *Consider the* 3*-graphs* K *and* L *in Figures* VII.4*a and b respectively. Let* $C_1 = \{b_5, c_8, a_7, b_8, a_8, c_6, a_6\}$, $C_2 = \{b_6, c_7, a_7, b_7, c_6\}$, $C_3 = \{b_1, a_2, c_3, a_4, b_4, c_2, a_1\}$ *and* $C_4 = \{b_2, c_3, b_3, a_3, c_2\}$. *Let* $\boldsymbol{S} = \{C_1, C_2, C_3, C_4\}$. *One can easily check that* $\boldsymbol{S}$ *is a spanning semicycle cover. Let* B_1, B_2, B_3 *and* B_4 *be* *the red-yellow bigons in L that contain* a_1, a_2, a_6 *and* a_5 *respectively. Let Q be the partition* $\{\{B_1\}, \{B_2, B_3\}, \{B_4\}\}$ of $B(L)$. *It is clear that L is partially congruent to K with respect to the partition Q. Furthermore L is faithful to K with respect to S. We note that* B_2 *and* B_3 *belong to distinct components of L, in agreement with Lemma* VII.10. ❑

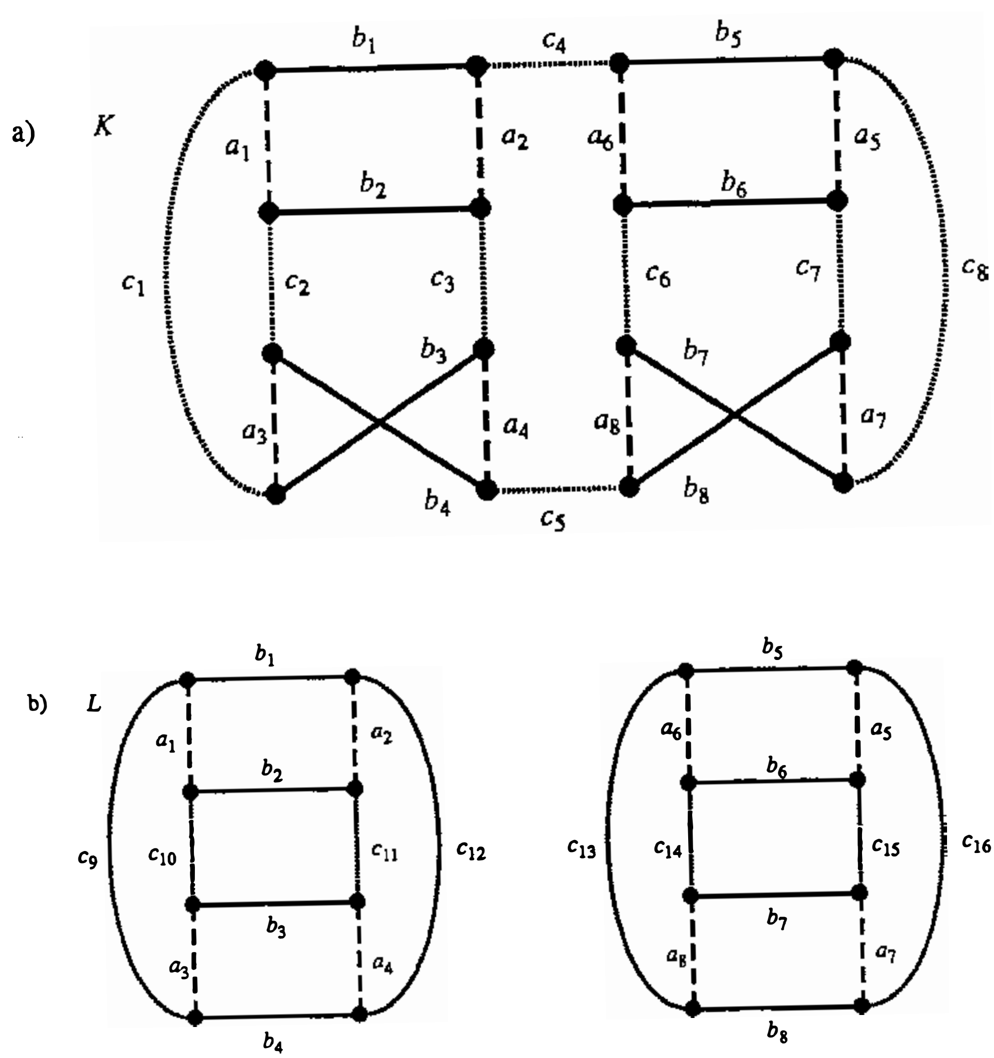

FIGURE VII.4

6. 3-GRAPH ENCODINGS

Let S be a semicycle cover in a 3-graph K. The semicycles of S induce a partition of VK into pairs, where two vertices belong to the same cell in this partition if and only if they are the two poles of a semicycle in S with respect to some red-yellow bigon. Let $\boldsymbol{V}$ denote this partition.

Let L be the 3-graph obtained from K by deleting the yellow edges and inserting a yellow edge joining the vertices v and w for each $\{v, w\} \in \boldsymbol{V}$. We say that L *encodes* K (with respect to S.) Clearly L is partially congruent to K with respect to a partition $\boldsymbol{Q}$. We say that $\boldsymbol{Q}$ is the *encoding partition*. Furthermore we have a one to one correspondence between the semicycles in S and the blue-yellow bigons in L. Hence L is faithful to K.

EXAMPLE VII.12. *Returning to Example* VII.11, *we see that L is in fact the* 3-*graph that encodes K with respect to S, and that $\boldsymbol{Q}$ is the encoding partition. We also note that L is planar. Lemma* VII.13 *below states that L will be planar in general whenever S is a spanning semicycle cover.* ❑

LEMMA VII.13. *Let L be the* 3-*graph that encodes a* 3-*graph K with respect to some spanning semicycle cover S. Then L is a planar* 3-*graph.*

Proof. Recall from Chapter III that L is planar if and only if the bigons of L span $\mathbf{Z}(L)$.

Let C be a circuit in L. By Lemma VII.9, $\beta(C)$ is the blue edge set of a cycle in K. Since $\boldsymbol{S}$ is spanning, then $\beta(C)$ is the blue edge set of a cycle $\Sigma(\boldsymbol{U} \cup \boldsymbol{Y})$ for a set $\boldsymbol{U}$ of semicycles in $\boldsymbol{S}$ and a set $\boldsymbol{Y}$ of red-blue bigons in K. Let $\boldsymbol{R}$ denote the set of blue-yellow bigons in L that correspond to the semicycles in $\boldsymbol{U}$. Evidently $\beta(C) = \beta(\Sigma(\boldsymbol{U} \cup \boldsymbol{Y})) = \beta(\Sigma(\boldsymbol{R} \cup \boldsymbol{Y}))$, and $\Sigma(\boldsymbol{R} \cup \boldsymbol{Y})$ is a cycle in L. By Lemma VII.3, $C = \Sigma(\boldsymbol{B} \cup \boldsymbol{R} \cup \boldsymbol{Y})$ for some set $\boldsymbol{B}$ of red-yellow bigons in L. Hence C is a sum of bigons, and we conclude that L is planar. ❑

7. COALESCING RED-YELLOW BIGONS

Let A and B be distinct red-yellow bigons of a 3-graph K. Let c_1 be a yellow edge in A and c_2 a yellow edge in B. Let c_1 join v_1 and w_1 and let c_2 join v_2 and w_2. Let L be the 3-graph obtained from K by deleting c_1 and c_2 and inserting two new yellow edges that join v_1 to v_2 and w_1 to w_2 respectively. We say that L is a 3-graph obtained by *coalescing* A and B. If A and B belong to distinct components of K, then clearly $c(L) = c(K) - 1$.

The following lemma is immediate.

LEMMA VII.14. *Suppose L is partially congruent to a 3-graph K with respect to a partition Q. Furthermore suppose J is obtained* *from L by coalescing two red-yellow bigons that belong to the same cell in Q. Then J is partially congruent to K.* ❑

LEMMA VII.15. *Let A and B be red-yellow bigons that belong to distinct components of a 3-graph K. Let L denote a 3-graph obtained from K by coalescing A and B. Then $\chi(L) = \chi(K) - 2$. Hence L is planar if and only if K is planar.*

Proof. Clearly $|VL| = |VK|$. However the number of red-yellow bigons has dropped by 1 as has the number of blue-yellow bigons. Hence

$$\begin{aligned}\chi(L) &= |B(K)| + |R(K)| + |Y(K)| - 2 - \frac{|VK|}{2} \\ &= \chi(K) - 2. \ \square\end{aligned}$$

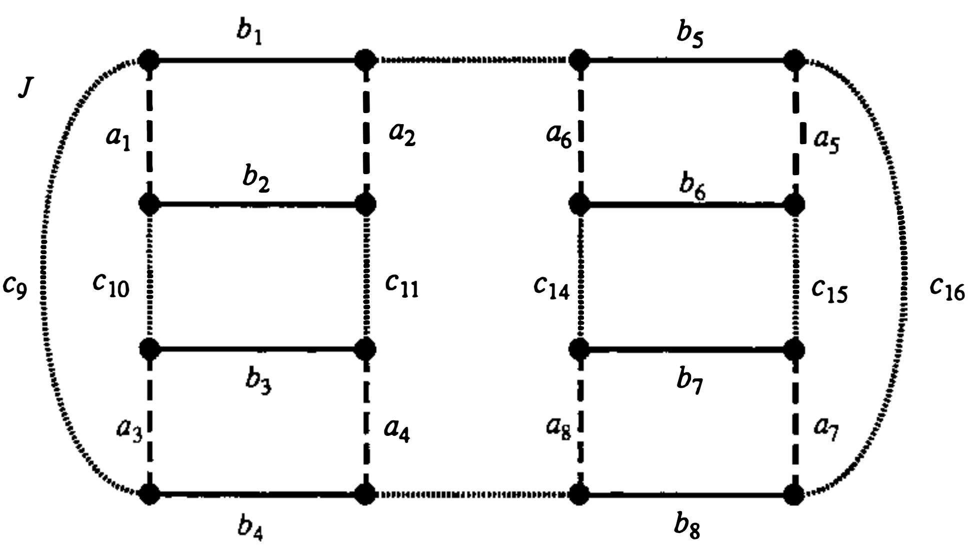

FIGURE VII.5

EXAMPLE VII.16. *Consider the* 3-*graphs* K *and* L *of Example* VII.11. *The* 3-*graph* J *in Figure* VII.5 *below is a* 3-*graph obtained from* L *by coalescing* B_2 *and* B_3. *We note that* J *is a planar* 3-*graph congruent to* K. ❑

THEOREM VII.17. *If there exists a spanning semicycle cover* S *in a* 3-*graph* K, *then* K *is congruent to a planar* 3-*graph.*

Proof. Let L_1 be the 3-graph that encodes K with respect to S. Let Q_1 be the encoding partition. By Lemma VII.13 L_1 is planar. Recall that L_1 is partially congruent to K with respect to Q_1 and L_1 is faithful to K. If each cell of Q_1 is a singleton, then L_1 is congruent to K and we are done.

Now suppose that there are two distinct red-yellow bigons B_1 and B_2 that belong to a common cell of Q_1. Since S is a spanning semicycle cover, then by Lemma VII.10 B_1 and B_2 belong to distinct components of L_1. Let L_2 denote the 3-graph obtained from L_1 by coalescing B_1 and B_2. By Lemma VII.15, L_2 is planar. By Lemma VII.14, L_2 is partially congruent to K with respect to a partition Q_2. Furthermore, for any $C \in S$, $|\beta(C) \cap \partial VB_1|$ and $|\beta(C) \cap \partial VB_2|$ are both even, and therefore $|\beta(C) \cap \partial VB|$ is even for all $B \in B(L_2)$. Hence we conclude that L_2 is faithful to K.

Proceeding inductively, we obtain a 3-graph L_n, partially congruent to K with respect to a partition Q_n, such that each cell in

Q_n is a singleton. Hence L_n is congruent to K. By repeated applications of Lemma VII.15, L_n is planar. ❑

8. THE EXISTENCE OF SPANNING SEMICYCLE COVERS

In this section we prove the converse of Theorem VII.17. This converse appears as Theorem VII.19, and it states that any 3-graph congruent to a planar 3-graph has a spanning semicycle cover.

LEMMA VII.18. *Let G be the red-yellow reduction of a 3-graph K. Let $\boldsymbol{C} = (C_1, C_2, ..., C_n)$ be a family of circuits in G and let $\boldsymbol{D} = \{D_1, D_2, ..., D_n\}$ be a set of semicycles in K such that $\beta(D_i) = C_i$ for all i. Then $\boldsymbol{C}$ spans $\boldsymbol{Z}(G)$ if and only if $\boldsymbol{B}(K) \cup \boldsymbol{D}$ spans $\boldsymbol{Z}(K)$.*

Proof. Firstly, assume that $\boldsymbol{C}$ spans $\boldsymbol{Z}(G)$ and let D be a circuit in K. We may assume that D is a b-cycle for otherwise D is a red-yellow bigon. Evidently $\beta(D)$ is a cycle in G and therefore $\beta(D) = \sum \boldsymbol{U}$ for some subset $\boldsymbol{U}$ of $\boldsymbol{C}$. Let $\boldsymbol{V}$ be the set of semicycles in $\boldsymbol{D}$ that correspond to the circuits in $\boldsymbol{U}$. Then clearly $\beta(\sum \boldsymbol{V}) = \beta(D)$. By Lemma VII.3, $D = \sum \boldsymbol{V} + \sum \boldsymbol{B}$ for some set $\boldsymbol{B} \subseteq \boldsymbol{B}(K)$. We conclude that $\boldsymbol{B}(K) \cup \boldsymbol{D}$ spans $\boldsymbol{Z}(K)$.

Now, assume that $\boldsymbol{B}(K) \cup \boldsymbol{D}$ spans $\boldsymbol{Z}(K)$ and let C be a circuit in G. Let D be a semicycle in K such that $\beta(D) = C$. Then $D = \sum \boldsymbol{V}$ for some subset $\boldsymbol{V}$ of $\boldsymbol{B}(K) \cup \boldsymbol{D}$. Let $\boldsymbol{U}$ be the set of circuits in $\boldsymbol{C}$ that correspond to the semicycles in $\boldsymbol{V} \cap \boldsymbol{D}$. Evidently, $C = \beta(D) =$

$\beta(\sum V) = \beta(\sum(V \cap D)) = \sum U$. We conclude that C spans $Z(G)$, as required. ❑

THEOREM VII.19. *If K is congruent to a planar 3-graph then there exists a spanning semicycle cover in K.*

Proof. Let L be a planar 3-graph congruent to K, and let G be the red-yellow reduction of K and L. By Theorem VII.6, there exists a semicycle cover S' of L that spans the boundary space of L. The fact that L is planar, S' spans the cycle space of L, and hence is a spanning semicycle cover of L. Let Y be the set of cycles in G that correspond to the red-blue bigons in K. Let R be the set of circuits in G that correspond to the semicycles in S'. By Lemma VII.18, we conclude that $Y \cup R$ spans $Z(G)$.

Let C_i denote a semicycle in K that represents D_i for each $D_i \in R$, and let S be the set of all such semicycles. The fact that R is a partition of EG establishes that the set of intersections of members of S with $\beta(K)$ is a partition of $\beta(K)$. Let D be a circuit in K. Since an even number of blue edges in D are incident on a given red-yellow bigon, then $\beta(D)$ is a cycle in G. Therefore $\beta(D) = \sum U$ for some set U consisting of cycles in $Y \cup R$. Let V be the set of semicycles in S that represent the circuits in $U \cap R$, and let W be the set of red-blue bigons in K that represent the circuits in $U \cap Y$. Therefore $\beta(D) = \sum U = \beta(\sum(V \cup W))$. Hence, by Lemma VII.3, $D = \sum(V \cup W) + \bigcup B$, for some set B of red-yellow bigons in $N(D)$.

Thus D is a sum of bigons in $B(K) \cup Y(K) \cup S$. We conclude that S is a spanning semicycle cover in K, as required. ❑

9. MACLANE'S THEOREM

The following theorem follows from Theorems VII.17 and VII.19.

THEOREM VII.20. *A 3-graph K is congruent to a planar 3-graph if and only if there exists a spanning semicycle cover in K.*

In this section we specialise Theorem VII.20 to the case of gems to obtain MacLane's theorem. If C is a cycle in a 3-graph K, then we denote by $N_Y(C)$ the set of all red-blue bigons that meet $\beta(C)$.

THEOREM VII.21 [MACLANE]. *A graph is planar if and only if its foundation has a spanning circuit double cover.*

Proof. Firstly, let G be a planar graph. Therefore G underlies a planar gem K. By Theorem VII.20 there exists a spanning semicycle cover S in K. Suppose there is a semicycle $S \in S$ that contains both blue edges of a red-blue bisquare Y. Therefore S corresponds to a path $\{e\}$ of length 1 in G. Assume that e is not an isthmus and hence that there exists a circuit C in G that contains e. Let D denote a semicycle in K that represents C. Then D contains only one blue edge of Y. Since no circuit in $B(K) \cup Y(K) \cup S$ contains just one blue edge of Y we have a contradiction to the fact that $B(K) \cup Y(K) \cup S$ spans $Z(K)$. Hence we conclude that each semicycle of S represents a circuit or corresponds to an isthmus in G. Let $S' = \{S_1, S_2, ..., S_m\}$ denote the set of semicycles in S that represent circuits in G. Clearly a semicycle in $S - S'$ is the sum of a red-blue bisquare and red-yellow bigons. Hence $B(K) \cup Y(K) \cup S'$ spans the cycle space of K.

Let C_i be the circuit of G that corresponds to S_i for each i. Since the set of intersections of members of S with $\beta(K)$ is a partition of $\beta(K)$, each red-blue bigon belongs to $N_Y(S_i)$ for exactly two values of i. It then follows that each edge in G that is not an isthmus belongs to C_i for exactly two values of i. Hence $C = (C_1, C_2, ..., C_n)$ is a circuit double cover for the foundation of G. We are required to show that $C_1, C_2, ..., C_n$ span $Z(G)$. Let D be a circuit in G. Let S be a semicycle in K that represents D. Hence $S = \sum U$ for a set U of circuits in $B(K) \cup Y(K) \cup S'$. Evidently no red-blue bigon has both of its blue edges in S, for otherwise S would not represent a circuit in G. Therefore $S + \sum U_1$ is also a semicycle that represents D, where U_1 is a set of red-yellow bigons in $N(S)$ and red-blue bigons in $N_Y(S)$. Let U_2 be the set of red-blue bigons in U that are not in $N_Y(S)$. Then $\sum U_2 + \sum(S' \cap U) + \bigcup B$ is also a semicycle that represents D, where B is the set of red-yellow bigons included in $\sum U_2 + \sum(S' \cap U)$. Let V be the set of components of C that correspond to the semicycles in $S' \cap U$. Then it follows that $\sum V = D$ since the edges of G that appear in exactly two circuits in V correspond to the red-blue bigons in U_2. We conclude that the components of $\boldsymbol{C}$ span $\boldsymbol{Z}(G)$, as required.

Now suppose the foundation of G has a spanning circuit double cover $(C_1, C_2, ..., C_m)$. Let K be a gem that G underlies. Let S_i be a semicycle in K that represents C_i, and let $\boldsymbol{S} = \{S_1, S_2, ..., S_m\} \cup \{S_{m+1}, S_{m+2}, ..., S_n\}$, where $\{S_{m+1}, S_{m+2}, ..., S_n\}$ is the set of red-blue bigons in K that correspond to the isthmuses in G. Since $S_i + Y$ is a semicycle that also represents C_i, where $Y \in N_Y(S_i)$, and each red-blue bigon not in $\{S_{m+1}, S_{m+2}, ..., S_n\}$ belongs to $N_Y(S_i)$ for exactly two values of i, we may choose each S_i so that the set of intersections of members of $\boldsymbol{S}$ with $\beta(K)$ is a partition of $\beta(K)$.

We now show that $\boldsymbol{B}(K) \cup \boldsymbol{Y}(K) \cup \boldsymbol{S}$ spans $\boldsymbol{Z}(K)$. Let D be a cycle in K. Suppose Y is a red-blue bigon such that both blue edges of Y belong to D. Then $D + Y$ is a sum of circuits in $\boldsymbol{B}(K) \cup \boldsymbol{Y}(K) \cup \boldsymbol{S}$ if and only if D is a sum of circuits in $\boldsymbol{B}(K) \cup \boldsymbol{Y}(K) \cup \boldsymbol{S}$. We therefore assume that D does not contain both blue edges of a red-blue bigon. Since D is a cycle, the number of blue edges incident on a given red-yellow bigon is always even. Therefore the set of edges that correspond to the bigons of $N_Y(D)$ is a cycle C in G. Since $C_1, C_2, ..., C_m$ span $\boldsymbol{Z}(G)$ then $C = \sum V$ for some set $V \subseteq \{C_1, C_2, ..., C_m\}$. Let $\boldsymbol{U}$ be the set of semicycles in $\boldsymbol{S}$ that represent the circuits of V. Let $D_1 = \sum U$. Let $D_2 = D_1 + \sum U_3$ where U_3 is the set of red-blue bigons that have both blue edges in D_1.

Evidently D_2 is a cycle such that $N_Y(D_2) = N_Y(D)$ and is a sum of circuits in $B(K) \cup Y(K) \cup S$. By adding to D_2 bigons in $N_Y(D_2)$ we can obtain a cycle D_3 such that $\beta(D_3) = \beta(D)$. By Lemma VII.3, the fact that D_3 is a sum of circuits in $B(K) \cup Y(K) \cup S$ implies that D is also a sum of circuits in $B(K) \cup Y(K) \cup S$. We conclude that S is a spanning semicycle cover in K. By Theorem III.12 K is congruent to a planar gem K'. Evidently G underlies K' and hence G is planar. ❑

Chapter VIII

THE HOMOLOGY OF N-GRAPHS

1. INTRODUCTION

A *map* is classically described as a cellular decomposition of a surface. In [3, 26, 27] the approach to maps is by way of graph imbedding schemes. In [29], Vince formulated, in terms of edge coloured graphs, a purely combinatorial generalisation of a map to higher dimensions. This generalisation he called a combinatorial map. In this chapter we introduce the concepts of the Euler characteristic and the homology spaces of a combinatorial map, and then show how they are related. In effect, we are generalising some of the work in Chapter III. This theory is analogous to the theory of homology in algebraic topology.

In [5], a related topic called a crystallisation is surveyed. In particular, Ferri and Gagliardi [4] introduce the concept of a crystallisation move. In the final portion of this chapter, we show that a crystallisation move on a combinatorial map does not disturb the Euler characteristic.

We write $\binom{n}{i}$ for the *binomial coefficient* $\frac{n!}{i!(n-i)!}$.

2. COMBINATORIAL MAPS

Let K be an n-regular graph where $n \geq 1$. A *proper edge colouring* of K is a colouring of the edges so that adjacent edges receive distinct colours. An *n-graph* is defined as an ordered triple (K, P, O) where K is an n-regular graph endowed with a proper edge colouring P in n colours and O is a ordering of the n colours. We write $K = (K, P, O)$ when no ambiguity results. We denote by $I_n = \{1, 2, ..., n\}$, the set of n colours. Hence P is a map from EK onto I_n. For convenience, we say that a 0-graph is the graph with one vertex and no edges.

A *combinatorial map* is an n-graph for some $n \geq 0$.

Let K be an n-graph, and suppose R is a m-graph obtained from K by deleting edges and isolated vertices, where $m > 0$. Then R is an *m-subgraph* (of K). An *m-residue* (of K) is a connected m-subgraph of K. Hence the edge set of a 1-residue of K consists of a single edge in EK.

A 0-*residue* of K is defined to be a vertex of K. Therefore, the set of all 0-residues in K is VK. The edge set of a 2-residue is a circuit whose edges alternate between two colours, and hence we sometimes refer to the edge set of a 2-residue as a *bigon* (or *bicoloured polygon.*)

A *residue* is an m-residue for some m. If K is a n-graph and R is an n-residue of K then clearly R is a component of K. If R is an m-residue of K, where $m < n$, then we say that R is a *proper residue* of K. However, in both cases we write $R \prec K$ or $K \succ R$. We note that $R \prec R$ for any residue R.

We usually write $[e_1, e_2, ..., e_j]$ for a residue with edge set $\{e_1, e_2, ..., e_j\}$. An m-residue is the *trivial* m-residue if it has just two vertices (and hence m edges.)

EXAMPLE VIII.1. *Consider the* 4-*graph K of Figure* VIII.1. *Since K is connected, there is only one* 4-*residue of K, namely K itself. The* 3-*residues in K are* $[a_1, b_1, c_1]$, $[a_2, b_2, c_2]$, $[a_1, b_1, d_1, d_2, a_2, b_2]$, $[b_1, c_1, d_1, d_2, b_2, c_2]$ *and* $[a_1, c_1, d_1, d_2, a_2, c_2]$. *There are three bigons of length four and six bigons of length two in K.* ❑

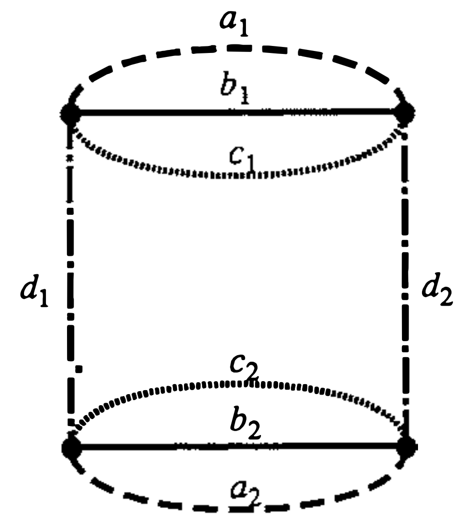

FIGURE VIII.1

3. THE SPACE OF m-CHAINS

We denote by $\boldsymbol{R}^m = \{R_1{}^m, R_2{}^m, ..., R_j{}^m\}$ the set of all m-residues of K. For $0 \leq m \leq n$, the *space* $\boldsymbol{C}_m(K)$ *of m-chains* of K is the vector space generated by the m-residues of K. Thus every m-chain of $\boldsymbol{C}_m(K)$ is a finite formal sum of the form $\sum_i z_i R_i^m$, where the coefficients z_i are elements of the field $GF(2)$ and $R_i^m \in \boldsymbol{R}^m$. Addition in this space satisfies the following rule. Let $C_1 = \sum_i y_i R_i^m$ and $C_2 = \sum_i z_i R_i^m$. Then $C_1 + C_2$ is the m-chain $\sum_i (y_i + z_i) R_i^m$.

If $C = \sum_i z_i R_i^m$ is an m-chain such that $z_i = 0$ for all i except for one value j of i, then C is the m-residue R_j^m.

The m-chain $\sum_i 0\ R_i^m$ is called the *zero m-chain* and we denote it by 0. For values of m other than 0, 1, ..., n we define $\boldsymbol{C}_m(K) = 0$.

4. THE BOUNDARY MAP

Let R be an m-residue of K for some integer m such that $0 < m \leq n$. The *boundary* $\delta_m R$ of R is the $(m-1)$-chain $\sum \boldsymbol{R}$, where $\boldsymbol{R}$ is the set of all $(m-1)$-residues of R. The boundary of a 0-residue is defined to be 0. The *boundary map* δ_m: $\boldsymbol{C}_m(K) \rightarrow \boldsymbol{C}_{m-1}(K)$ is the linear transformation defined for $0 \leq m \leq n$ by $\delta_m \sum_i z_i R_i^m = \sum_i z_i \delta_m R_i^m$ and defined to be trivial otherwise. We sometimes write δ for δ_m when no ambiguity results.

EXAMPLE VIII.2. *Consider the* 4*-graph of Example* VIII.1. *Let* C *be the* 3*-chain* $[a_1, b_1, c_1] + [a_1, b_1, d_1, d_2, a_2, b_2]$. *The boundary* δC *of* C *is then the* 2*-chain* $[a_1, b_1] + [b_1, c_1] + [a_1, c_1] + [a_1, a_2, d_1, d_2] + [a_1, b_1] + [a_2, b_2] + [b_1, b_2, d_1, d_2]$. *The boundary* $\delta(\delta C)$ *of* δC *is the* 1*-chain* $[a_1] + [b_1] + [b_1] + [c_1] + [a_1] + [c_1] + [a_1] + [a_2] + [d_1] + [d_2] + [a_1] + [a_2] + [b_1] + [b_2] + [b_1] + [b_2] + [d_1] + [d_2] = 0$. *In general, Theorem* VIII.4 *below shows that* $\delta(\delta C) = 0$ *for any* m*-chain* C. ❑

5. M-CYCLES

The *kernel* of the boundary map δ_m consists of those m-chains with empty boundary. The elements of the kernel are *m-cycles*. Since δ_m is a linear transformation, the kernel of δ_m is a subspace of $\boldsymbol{C}_m(K)$ and we denote it by $\boldsymbol{Z}_m(K)$.

EXAMPLE VIII.3. *The* 2-*chain* $C = [a_1, b_1] + [b_1, c_1] + [c_1, a_1]$ *in the* 4-*graph of Figure* VIII.1 *is a* 2-*cycle since its boundary is* $[a_1] + [b_1] + [b_1] + [c_1] + [c_1] + [a_1] = 0$. *We note that C is the boundary of the* 3-*residue* $[a_1, b_1, c_1]$. *In fact, we shall see in general that all boundaries of m-residues are in fact* $(m - 1)$-*cycles.* ❑

6. M-BOUNDARIES

The *image* of the boundary map δ_{m+1} consists of those m-chains which are boundaries of $(m + 1)$-chains. The elements of the image are *m-boundaries.* Since δ_{m+1} is a linear transformation, the image of δ_{m+1} is a subspace of $\boldsymbol{C}_m(K)$ and we denote it by $\boldsymbol{B}_m(K)$. We note that $\dim \boldsymbol{B}_{-1}(K) = \dim \boldsymbol{B}_n(K) = 0$. In the next section we shall show that $\boldsymbol{B}_m(K)$ is in fact a subspace of $\boldsymbol{Z}_m(K)$.

7. $\delta_{m-1}\delta_m$ IS THE TRIVIAL MAP

If $I \subseteq I_n$ and R is a $|I|$-residue such that $P(ER) = I$ then we say that R is an *I-residue.* We now come to an equation that occurs analogously in many different branches of mathematics.

THEOREM VIII.4. *For every n-graph K the composite function* $\delta_{m-1}\delta_m$ *mapping* $\boldsymbol{C}_m(K)$ *into* $\boldsymbol{C}_{m-2}(K)$ *is* 0.

Proof. Since a linear transformation is completely determined by its values on its basis elements, it is enough to check that for an m-residue R, we have $\delta_{m-1}\delta_m(R) = 0$.

Since δ_m is defined to be 0 for values of m other than 1, 2, ..., n, we consider just the following two cases.

Case i) $m = 2$. Let $R = [e_1, e_2, ..., e_r]$ be a 2-residue where e_i is adjacent to e_{i+1} for each $i < r$. (It follows that e_r is adjacent to e_1.) Let v_i be the vertex which is incident on e_i and e_{i+1} and let v_r be the vertex incident on e_r and e_1. Then $\delta R = [e_1] + [e_2] + ... + [e_r]$. This implies that $\delta(\delta R) = \delta([e_1] + [e_2] + ... + [e_r]) = (v_r + v_1) + (v_1 + v_2) + (v_2 + v_3) + ... + (v_{r-1} + v_r) = 0$ since each v_i appears twice in this sum. Hence $\delta_1\delta_2 = 0$.

Case ii) $m \geq 3$. Let R be an m-residue in K. Let R' be an $(m - 2)$-residue that is a residue of R. Hence $\boldsymbol{P}(ER') = \boldsymbol{P}(ER) - \{i, j\}$ where $i, j \in \boldsymbol{P}(ER)$, and $i \neq j$. Clearly R' is a residue of a unique $(m - 1)$-residue R_1 where $\boldsymbol{P}(ER_1) = \boldsymbol{P}(ER') \cup \{i\}$, and a residue of a unique $(m - 1)$-residue R_2 where $\boldsymbol{P}(ER_2) = \boldsymbol{P}(ER') \cup \{j\}$. Hence we have $R' \prec R_1 \prec R$ and $R' \prec R_2 \prec R$. Therefore the coefficient of R' in the sum $\delta(\delta R)$ is $2 \equiv 0 (\text{mod } 2)$. Since R' is any $(m - 2)$-residue of R, it follows that $\delta(\delta R) = 0$, and hence $\delta_{m-1}\delta_m = 0$. ❑

COROLLARY VIII.5. $\boldsymbol{B}_m(K)$ *is a subspace of* $\boldsymbol{Z}_m(K)$.

Proof. By definition $B_m(K) = \delta_{m+1}(C_{m+1}(K))$. If $B \in B_m(K)$ then $B = \delta_{m+1}(C)$ for some $C \in C_{m+1}(K)$. Thus $\delta_m(B) = \delta_m(\delta_{m+1}(C)) = 0$, and hence $B \in Z_m(K)$. ❑

The orthogonal complement $H_m(K)$ of $B_m(K)$ in $Z_m(K)$ is the mth *homology space* of K. The dimension of $H_m(K)$ is called the mth *connectivity number* of K, and is denoted by $h_m(K) = h_m$. Hence $h_m(K) = \dim Z_m(K) - \dim B_m(K)$.

8. THE EULER CHARACTERISTIC OF AN n-GRAPH

Let K be an n-graph. If $0 \le m \le n$, we define $r_m(K) = r_m = |R^m|$, the number of m-residues in K. Hence $r_0(K) = |VK|$, $r_1(K) = |EK|$ and $r_n(K) = c(K)$. Furthermore, if K is a 3-graph then $r_2(K) = r(K)$.

The *Euler characteristic* of an n-graph K is defined to be $\chi'(K) = \sum_{i=0}^{n}(-1)^i r_i(K)$.

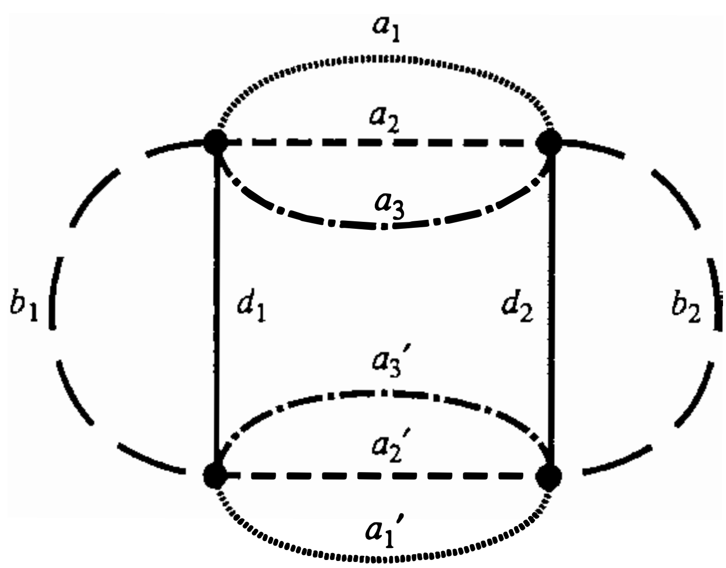

FIGURE VIII.2

EXAMPLE VIII.6. *Consider the 5-graph K of Figure* VIII.2. *One may easily check that* $r_0(K) = 4$, $r_1(K) = 10$, $r_2(K) = 14$, $r_3(K) = 11$, $r_4(K) = 5$ *and* $r_5(K) = 1$. *Hence the Euler characteristic of K is* $\chi'(K) = 4 - 10 + 14 - 11 + 5 - 1 = 1$. ❑

We are now in a position to state and prove our main theorem for this chapter. Theorem VIII.7 below relates the Euler characteristic to a function of the h_i's.

THEOREM VIII.7. *For any n-graph K,*

$$\chi'(K) = \sum_{i=0}^{n} (-1)^i h_i(K).$$

Proof. The dimension of the domain of δ_i is $r_i = |\boldsymbol{R}^i|$. The dimension of the image space of δ_i is $\dim \boldsymbol{B}_{i-1}(K)$ and the dimension of the kernel of δ_i is $\dim \boldsymbol{Z}_i(K)$. By linear algebra, we have that $r_i = \dim \boldsymbol{B}_{i-1}(K) + \dim \boldsymbol{Z}_i(K)$. Therefore $r_i - h_i = \dim \boldsymbol{B}_{i-1}(K) + \dim \boldsymbol{Z}_i(K) - \dim \boldsymbol{Z}_i(K) + \dim \boldsymbol{B}_i(K) = \dim \boldsymbol{B}_{i-1}(K) + \dim \boldsymbol{B}_i(K)$. Thus $\sum_{i=0}^{n} (-1)^i (r_i - h_i) = \dim \boldsymbol{B}_{-1}(K) \pm \dim \boldsymbol{B}_n(K) = 0$, since $\dim \boldsymbol{B}_{-1}(K) = \dim \boldsymbol{B}_n(K) = 0$. Hence $\chi'(K) = \sum_{i=0}^{n} (-1)^i r_i = \sum_{i=0}^{n} (-1)^i h_i$, as required. ❑

9. DIPOLES

Let v and w be a pair of adjacent vertices in an n-graph K. Suppose that v and w are joined by exactly m edges $e_1, e_2, \ldots, e_m$, where $0 < m < n$, and let R be the trivial m-residue $[e_1, e_2, \ldots, e_m]$.

Let A and B denote the $(P(EK) - P(ER))$-residues that contain v and w respectively. We say that R is an *m-dipole* if A and B are distinct residues.

EXAMPLE VIII.8. *Consider the* 5*-graph* K *in Figure* VIII.2. *Then the* 3*-dipoles in* K *are* $[a_1, a_2, a_3]$ *and* $[a_1', a_2', a_3']$. *The* 2*-dipoles in* K *are* $[b_1, d_1]$ *and* $[b_2, d_2]$. *The are no* 1*-dipoles or* 4*-dipoles in* K. ❑

10. CANCELLATIONS AND CREATIONS OF DIPOLES

Let R be an m-dipole in a n-graph K, and let $VR = \{v, w\}$. Let a_i and a_i' be the edges coloured i incident on v and w respectively, where $i \in P(EK) - P(ER)$. Let v_i, w_i be the vertices other than v and w incident on a_i and a_i' respectively. The *cancellation* of this m-dipole R is the operation of deletion of the vertices v and w followed by the insertion of an edge joining v_i to w_i for each $i \in P(EK) - P(ER)$. We denote the resulting graph by $K - [R]$. We observe that A and B have coalesced into one $(P(EK) - P(ER))$-residue. The *creation* of an m-dipole is the inverse operation.

EXAMPLE VIII.9. *Figure* VIII.3 *illustrates the resulting* 5*-graph* K' *obtained by cancelling the* 2*-dipole* $[b_2, d_2]$ *in the* 5*-graph* K *of Figure* VIII.2. *We note that* $r_0(K') = 2$, $r_1(K') = 5$, $r_2(K') = \binom{5}{2} = 10$, $r_3(K') = \binom{5}{3} = 10$, $r_4(K') = \binom{5}{4} = 5$ *and* $r_5(K') = 1$. *Hence* $\chi'(K') = 2 - 5 + 10 - 10 + 5 - 1 = 1 = \chi'(K)$. *Theorem* VIII.17 *below* *shows that, in general, cancellation of an m-dipole does not alter the residue characteristic.* ❑

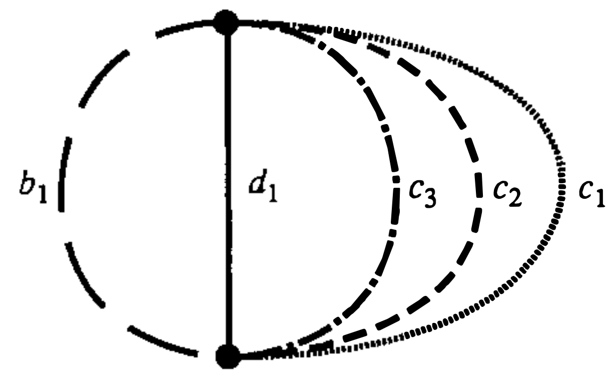

FIGURE VIII.3

11. BALANCED DIPOLES

Let R be an m-dipole in an n-graph K, and let $VR = \{v, w\}$. Let Q be an i-residue of K where $0 < i \leq n$. Consider the following cases.

i) If $EQ \cap ER \neq \varnothing$ then we say that Q *crosses* R.

ii) If $EQ \cap ER = \varnothing$ and $\{v, w\} \cap VQ = \varnothing$ then we say that Q *avoids* R.

iii) If $EQ \cap ER = \varnothing$ and $v \in VQ$ then we say that Q *kisses* R *(at v)*.

iv) If $EQ \cap ER = \varnothing$ and $w \in VQ$ then we say that Q *kisses* R *(at w)*.

LEMMA VIII.10. *Let R be an m-dipole in an n-graph K and let $VR = \{v, w\}$. Then no i-residue kisses R at v and w.*

Proof. Suppose Q is an i-residue that kisses R at v and w. Let A and A' be the $(P(EK) - P(ER))$-residues that contain v and w respectively. Since Q kisses R at v, then $Q \prec A$. Furthermore, since Q kisses R at w, then $Q \prec A'$. Therefore $A = A'$, contradicting the fact that R is an m-dipole. ❑

COROLLARY VIII.11. *Let R be an m-dipole in an n-graph K and let $VR = \{v, w\}$. If Q is an i-residue in K where $0 < i \leq n$, then either $Q \prec R$, Q crosses R, Q avoids R, Q kisses R at v, or Q kisses R at w.* ❑

Suppose Q is an i-residue that kisses an m-dipole R at a vertex v. Let w be the vertex other than v in VR. Let Q' be the unique $(P(EQ))$-residue that kisses R at w. We say that Q' is the *mate* of Q (with respect to R) and that the i-subgraph of K with components Q and Q' *hugs* R.

EXAMPLE VIII.12. *Consider the 5-graph K of Figure VIII.2 and let R be the 2-dipole $[b_2, d_2]$. $Q = [a_1, a_2, d_1, d_2, a_1', a_2']$ is an example of a 3-residue that crosses R. The 2-subgraph Q' with components $[a_1, a_2]$ and $[a_1', a_2']$ is an example of a 2-graph that hugs R.* ❑

Let R be an m-dipole in an n-graph K. Suppose Q is an i-subgraph such that either

i) Q is an i-residue that avoids R or,

ii) Q is an i-residue that crosses R or,

iii) Q is an i-subgraph with two components that hugs R.

Then we say that Q *balances* R.

LEMMA VIII.13. *Let R be an m-dipole in an n-graph K. There is a one to one correspondence between the i-residues of $K - [R]$ and the i-subgraphs that balance R, whenever $0 < i \leq n$.*

Proof. Suppose Q is an i-subgraph that balances R, and let $VR = \{v, w\}$. Let $X = \{a_1, a_2, ..., a_j\} = (\partial v \cap EQ) - ER$ and let $W = \{b_1, b_2, ..., b_j\} = (\partial w \cap EQ) - ER$. Since $\boldsymbol{P}(X) = \boldsymbol{P}(W)$, we choose the labelling of W so that $\boldsymbol{P}(a_k) = \boldsymbol{P}(b_k)$ for all k. Let $\psi a_k = \{v, v_k\}$ and $\psi b_k = \{w, w_k\}$ for all k such that $1 \leq k \leq j$. In $K - [R]$ let c_k denote the edge coloured $\boldsymbol{P}(a_k) = \boldsymbol{P}(b_k)$ joining v_k and w_k. Let $U = \{c_1, c_2, ..., c_k\}$. Then clearly $(EQ - ER - X - W) \cup U$ is the edge set of a unique i-residue Q' in $K - [R]$ such that $\boldsymbol{P}(Q) = \boldsymbol{P}(Q')$. Thus we have a correspondence between the i-subgraphs of K that balance R and the i-residues of $K - [R]$.

By reversing the above argument, one can show that any i-residue in $K - [R]$ corresponds to a unique i-subgraph of K that balances R. ❑

LEMMA VIII.14. *Suppose R is an m-dipole in the n-graph K. Then the number of i-subgraphs that hug R is $\binom{n-m}{i}$, where $0 < i \leq n - m$.*

Proof. Let $VR = \{v, w\}$. The number of i-residues that contain i edges in $\partial v - ER$ is $\binom{n-m}{i}$. Let Q be such an i-residue and let Q' be the mate of Q with respect to R. Therefore Q and Q' are the components of a i-subgraph that hugs R. Hence there is a total of $\binom{n-m}{i}$ such i-subgraphs. ❑

EXAMPLE VIII.15. *Returning to Example* VIII.12, *we see that the number of* 2*-residues in* $[a_1, a_2, a_3]$ *is* $\binom{3}{2} = 3$ *which is the number of* 2*-residues in* $[a_1', a_2', a_3']$. *Hence* $[a_1, a_2, a_1', a_2']$, $[a_2, a_3, a_2', a_3']$ *and* $[a_1, a_3, a_1', a_3']$ *are the three* 2*-subgraphs that hug* R. *They correspond to the* 2*-residues* $[c_1, c_2]$, $[c_2, c_3]$, *and* $[c_1, c_3]$ *respectively in the* 5*-graph of Figure* VIII.3. ❑

We now have the following lemma.

LEMMA VIII.16. *Let* R *be an* m*-dipole of an* n*-graph* K, *and let* $a = min\{m, n - m\}$ *and* $b = max\{m, n - m\}$. *Then*

$$r_i(K - [R]) = \begin{cases} r_i(K) - \binom{a}{i} - \binom{b}{i} & \text{if } i \le a \\ r_i(K) - \binom{b}{i} & \text{if } a < i \le b \\ r_i(K) & \text{if } b < i \le n. \end{cases}$$

Proof. Let $VR = \{v, w\}$. Consider the following cases.

Case a) $i \le a$. If $i = 0$, then $r_0(K - [R]) = |VK - [R]| = |VK| - 2 = r_0(K) - 2 = r_0(K) - \binom{a}{0} - \binom{b}{0}$, as required. Now assume that

$0 < i$, and let Q be an i-residue of K. Since $i \leq m$ and $i \leq n - m$, then either $Q \prec R$, Q avoids R, Q crosses R, or Q together with its mate are the components of an i-subgraph that hugs R. The number of i-residues of R is $\binom{m}{i}$, and by Lemma VIII.14 the number of i-subgraphs that hug R is $\binom{n-m}{i}$. Since each i-subgraph that hugs R consists of two i-residues, then the number of i-subgraphs that balance R is

$$r_i(K) - \binom{m}{i} - \binom{n-m}{i} = r_i(K) - \binom{a}{i} - \binom{b}{i}.$$

Case b) $a < i \leq b$. Consider the following two subcases.

Subcase i) $a = m$ and $b = n - m$. Let Q be an i-residue of K. Since $m < i$ then Q cannot be an i-residue of R. Therefore either Q avoids R, Q crosses R, or Q together with its mate are the components of an i-subgraph hug R. Since each i-subgraph that hugs R consists of two i-residues, then the number of i-subgraph that balance R is

$$r_i(K) - \binom{n-m}{i} = r_i(K) - \binom{b}{i}.$$

Subcase ii) $a = n - m$ and $b = m$. Let Q be an i-residue of K. Since $n - m < i$ then Q cannot kiss R. Hence all i-subgraphs that balance R are i-residues and i-residues that do not balance R must be i-residues of R. Therefore the number of i-subgraphs that balance R is

$$r_i(K) - \binom{m}{i} = r_i(K) - \binom{b}{i}.$$

Case c) $b < i \leq n$. Let Q be an i-residue of K. Since $m < i$ and $n - m < i$ then Q cannot be an i-residue of R and cannot kiss R. Hence all i-residues in K must balance R. Therefore the number of i-subgraphs that balance R is $r_i(K)$.

The lemma now follows from Lemma VIII.13. ❑

The following theorem is a generalisation of Theorem II.3 in Chapter II.

THEOREM VIII.17. *If R is an m-dipole in an n-graph K then* $\chi'(K - [R]) = \chi'(K)$.

Proof. It is a well known fact that

$$\sum_{i=0}^{k} (-1)^i \binom{k}{i} = 0$$

if $k > 0$. Let $a = \min\{m, n - m\}$ and $b = \max\{m, n - m\}$. Then by Lemma VIII.16,

$$\chi'(K-[R]) = \sum_{i=0}^{n}(-1)^i r_i(K-[R])$$

$$= \sum_{i=0}^{a}(-1)^i r_i(K-[R]) + \sum_{i=a+1}^{b}(-1)^i r_i(K-[R])$$

$$+\sum_{i=b+1}^{n}(-1)^i r_i(K-[R])$$

$$= \sum_{i=0}^{a}(-1)^i\left(r_i(K) - \binom{a}{i} - \binom{b}{i}\right) + \sum_{i=a+1}^{b}(-1)^i\left(r_i(K) - \binom{b}{i}\right)$$

$$+\sum_{i=b+1}^{n}(-1)^i r_i(K)$$

$$= \sum_{i=0}^{n}(-1)^i r_i(K) - \sum_{i=0}^{a}(-1)^i\binom{a}{i} - \sum_{i=0}^{b}(-1)^i\binom{b}{i}$$

$$= \chi'(K). \square$$

REFERENCES

1. L. AUSLANDER, T.A. BROWN, AND J.W.T. YOUNGS, The imbedding of graphs in manifolds, *J. Math. and Mech.* **12** (1963), 629-634.

2. M. DEHN AND P. HEEGAARD, Analysis situs, *Encyklopädie der mathematischen Wissenschaften* **13** (1982), 295-308.

3. J.R. EDMONDS, A combinatorial representation for polyhedral surfaces, *Notices Amer. Math. Soc.* **7** (1960), 646.

4. M. FERRI AND C. GAGLIARDI, Crystallization moves, *Pacific J. Math* **100** (1982), 85-103.

5. M. FERRI, C. GAGLIARDI, AND L. GRASSELLI, A graph-theoretical representation of PL-manifolds – A survey on crystallizations, *Aequationes Mathematicae* **31** (1986), 121-141.

6. J.L. GROSS AND T.W. TUCKER, "Topological graph theory," Wiley-Interscience, New York (1987).

7. F. JAEGER, N.H. XUONG, AND C. PAYAN, Maximum genus and trees, *J. Graph Theory* **3** (1987), 387-391.

8. M. JUNGERMAN, Characterization of upper embeddable graphs, *Trans. Amer. Math. Soc.* **241** (1978), 401-406.

9. M.P. KHOMENKO, M.O. OSTROVERKI, AND V.O. KUZMENKO, Phi-transformations of graphs (in Russian). Institute of Mathematics of the Ukranian Academy of Sciences, pp. 384, 1973.

10. M.P. KHOMENKO, M.O. OSTROVERKI, AND V.O. KUZMENKO, Maximum subgraphs in phi-transformations of graphs (in Russian). Institute of Mathematics of the Ukranian Academy of Sciences, pp. 180-210, 1973.

11. S. LINS, "Graphs of maps," Ph.D. dissertation, University of Waterloo, 1980.

12. S. LINS, Graph-encoded maps, *J. Combin. Theory Ser. B* **32** (1982), 171-181.

13. C.H.C. LITTLE, Cubic combinatorial maps, *J. Combin. Theory Ser. B* **44** (1988), 44-63.

14. C.H.C. LITTLE AND A. VINCE, Embedding schemes and the Jordan curve theorem. In "Topics in Combinatorics and Graph Theory," R. Bodendiek and R. Henn, Physica-Verlag, Heidelberg, 1990, pp. 479-489.

15. L. NEBESKY, A new characterization of the maximum genus of a graph, *Czechoslovak Math. J.* **31** (1981), 604-613.

16. E.A. NORDHAUS, B.M. STEWART, AND A.T. WHITE, On the maximum genus of a graph, *J. Combin. Theory Ser. B* **11** (1971), 258-267.

17. E.A. NORDHAUS, R.D. RINGEISEN, B.M. STEWART, AND A.T. WHITE, A Kuratowski-type theorem for the maximum genus of a graph, *J. Combin. Theory Ser. B* **12** (1972), 260-267.

18. R.D. RINGEISEN, Upper and lower embeddable graphs. In "Graph Theory and Applications," Springer-Verlag, New York/Berlin, 1972.

19. R.D. RINGNEISEN, Survey of Results on the Maximum Genus of a Graph, *J. Graph Theory* **3** (1979), 1-13.

20. N. ROBERTSON, "Graphs minimal under Girth, Valency and Connectivity," Ph.D. dissertation, University of Waterloo, 1971.

21. M. SKOVIERA, Decay number and the maximum genus of a graph, Submitted for publication.

22. M. SKOVIERA, The maximum genus of graphs of diameter two, *Discrete Mathematics* **87** (1991), 175-180.

23. S. STAHL, Permutation-partition pairs: A combinatorial generalization of graph embedding, *Trans. Amer. Math. Soc.* **1** (1980), 129-145.

24. S. STAHL, A combinatorial analog of the Jordan curve theorem, *J. Combin. Theory Ser. B* **35** (1983), 28-38.

25. S. STAHL, A Duality for Permutations, *Discrete Math.* **71** (1988), 257-271.

26. W.T. TUTTE, Combinatorial oriented maps, *Canad. J. Math.* **5** (1979), 986-1004.

27. W.T. TUTTE, "Graph Theory," Addison-Wesley, Menlo Park, CA (1984).

28. A. VINCE, "Combinatorial Maps," Ph.D. dissertation, University of Michigan, Ann Arbor, Mich., 1980.

29. A. VINCE, Combinatorial Maps, *J. Combin. Theory Ser. B* **34** (1983), 1—21.

30. A. VINCE, Recognizing the 3-sphere. In "Graph Theory and its Applications: East and West," M.F. Capobianco, M. Guan, D.F. Hsu, and F. Tian, New York Academy of Sciences, December 1989, pp. 571-583.

31. A. VINCE AND C.H.C. LITTLE, Discrete Jordan Curve Theorems, *J. Combin. Theory Ser. B* **47** (1989), 251-261.

32. N.H. XUONG, How to Determine the Maximum Genus of a Graph, *J. Combin. Theory Ser. B* **26** (1979), 217-225.

33. N.H. XUONG, Upper-embeddable Graphs and Related Topics, *J. Combin. Theory Ser. B* **26** (1979), 226-232.

INDEX

Massey University Library

Thesis Copyright Form

Title of thesis: Combinatorial Maps and the Foundations of Topological Graph Theory.

(1) (a) I give permission for my thesis to be made available to readers in the Massey University Library under conditions determined by the Librarian.

(b) I do not wish my thesis to be made ~~available to readers without my written consent for ______ months.~~

(2) (a) I agree that my thesis, or a copy, may be sent to another institution under conditions determined by the Librarian.

(b) I do not wish my thesis, ~~or a copy, to be sent to another institution without my written consent for ______ months.~~

(3) (a) I agree that my thesis may be copied for Library use.

(b) I do not wish my thesis to be ~~copied for Library use for ______ months.~~

C.P. Bonnington

Signed Bonnington.

Date 2 Aug 1991

The copyright of this thesis belongs to the author. Readers must sign their name in the space below to show that they recognise this. They are asked to add their permanent address.

NAME AND ADDRESS	DATE
Bonnington C.P. BONNINGTON 5a WILFRED ST HAMILTON	2 AUG 1991