Map Graphs*

Zhi-Zhong Chen† Michelangelo Grigni‡ Christos H. Papadimitriou§

October 13, 1999

Abstract

We consider a modified notion of planarity, in which two nations of a map are considered adjacent when they share any *point* of their boundaries (not necessarily an *edge*, as planarity requires). Such adjacencies define a *map graph*. We give an NP characterization for such graphs, and an $O(n^3)$-time recognition algorithm for a restricted version: given a graph, decide whether it is realized by adjacencies in a map without holes, in which at most four nations meet at any point.

1 Introduction

1.1 Motivation and Definition

Suppose you are told that there are four planar regions, and for each pair you are told their topological relation: A is inside B, B overlaps C, C touches D on the outside, D overlaps B, D is disjoint from A, and C overlaps A. All four planar regions are "bubbles" with no holes; more precisely, they are closed disc homeomorphs. Are such regions possible? If so, we would like a model, a picture of four regions so related; if not, a proof of impossibility.

This extension of propositional logic is known as the *topological inference problem* [9]. No decision algorithm or finite axiomatization is known, although the problem becomes both finitely axiomatizable and polynomial-time decidable in any number of dimensions other than two [1, 14], for some reasonable vocabularies of topological relations. In fact, the following special case has been open since the 1960's [8]: for every pair of regions, we are only told whether the regions intersect or not. This is known as the *string graph problem*, because the input is a simple graph, and we may assume that the regions are in fact planar curves (or slightly fattened simple curves, if we insist on disc homeomorphs). In other words, we are seeking a recognition algorithm for the intersection graphs of planar curves. It is open whether this problem is decidable; it is known that there are infinitely many forbidden subgraphs, that recognition is at least NP-hard [11], and that there are string graphs that require exponentially many string intersections for their realization [12].

*A preliminary version appeared with the title "Planar Map Graphs" [4].

†Department of Mathematical Sciences, Tokyo Denki University, Hatoyama, Saitama 350-0394, JAPAN. Part of work done while visiting University of California at Berkeley. E-mail: chen@r.dendai.ac.jp.

‡Department of Mathematics and Computer Science, Emory University, Atlanta, GA 30322. E-mail: mic@mathcs.emory.edu.

§Computer Science Division, University of California at Berkeley, Berkeley, CA 94720. Supported by NSF Grant number CCR-9626361. E-mail: christos@cs.berkeley.edu.

The difficulty of the string graph problem stems to a large extent from the complex overlaps allowed between regions. But many practical applications are so structured that no two regions in them overlap arbitrarily. For example, consider maps of political regions: such regions either contain one another or else they have disjoint interiors; general overlaps are not allowed.

In this paper, we consider the following special case: for every pair of regions, they are either disjoint, or they intersect only on their boundaries. Since the nations of political maps intersect in this way, we call such regions *nations*. In this paper we study *map graphs*, the intersection graphs of nations.

Definition 1.1 *Suppose G is a simple graph. A* map *of G is a function $\mathcal{M}$ taking each vertex v of G to a closed disc homeomorph $\mathcal{M}(v)$ in the sphere, such that:*

1. *For every pair of distinct vertices u and v, the interiors of $\mathcal{M}(u)$ and $\mathcal{M}(v)$ are disjoint.*
2. *Two vertices u and v are adjacent in G if and only if the boundaries of $\mathcal{M}(u)$ and $\mathcal{M}(v)$ intersect.*

If G has a map, then it is a map graph. *The regions $\mathcal{M}(v)$ are the* nations *of $\mathcal{M}$. The uncovered points of the sphere fall into open connected regions; the closure of each such region is a* hole *of $\mathcal{M}$.*

Can we recognize map graphs? This problem is closely related to planarity, one of the most basic and influential concepts in graph theory. Usually, planarity is defined as above, but with adjacency only for those pairs of nations sharing a curve segment. Planarity may also be defined in terms of maps: specifically, a planar graph has a map *such that no four nations meet at a point.*

We consider two natural restrictions on maps and map graphs. First, suppose we restrict our map so that *no more than k nations meet at a point*; we call this a *k-map*, and the corresponding graph is a *k-map graph.* In particular the ordinary planar graphs are the 2-map graphs; in fact all 3-map graphs are also 2-map graphs, as argued below.

Second, suppose that every point of the sphere is covered by some nation. Then we say that this is a *hole-free map*, and the corresponding graph is a *hole-free map graph.* For our algorithm starting in Section 3, we consider hole-free 4-map graphs.

In our figures we draw a map by projecting one point of the sphere to infinity; we always choose a point that is not on a nation boundary.

1.2 Examples

We give three examples. First, consider the adjacency graph of the United States in Figure 1.1; this is a 4-map graph. It is not planar, since the "four corners" states (circled) form a K_4, which is part of a K_5-minor. Since removing one nation from a hole-free map leaves a connected set of nations, all hole-free map graphs are 2-connected; consequently, this example is not a hole-free map graph.

Second, consider the 17-nation hole-free 4-map in Figure 1.2(1). Let G be its map graph. At the end of Section 2.1, we show that after deleting the edge $\{6, 7\}$ from G, the

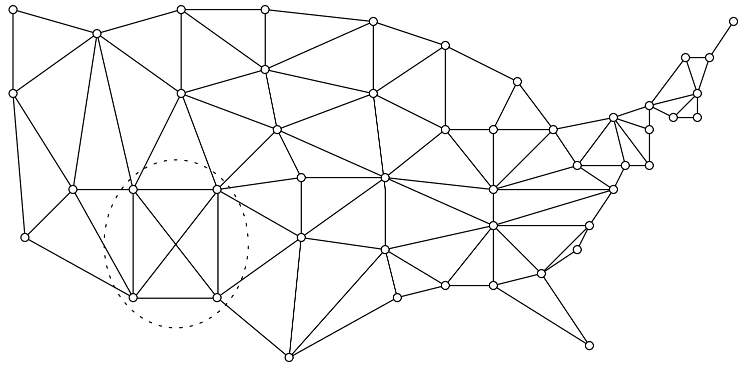

Figure 1.1: The USA map graph.

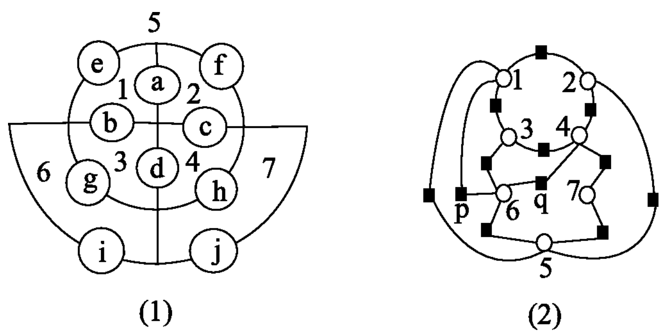

Figure 1.2: An example hole-free 4-map, and a subgraph of its witness.

result is not a map graph. This example demonstrates that the 4-map graph property is not monotone.

Third, consider Figure 1.3: part (1) is a map graph, part (2) is a hole-free 4-map of the graph, and part (3) is a corresponding witness (as defined in Section 2.1). The graph has a mirror symmetry exchanging a with c, but the map and the witness do not. In fact a careful analysis shows that no map or witness of this graph has such a symmetry, and so a layout algorithm must somehow "break symmetry" to find a map for this graph.

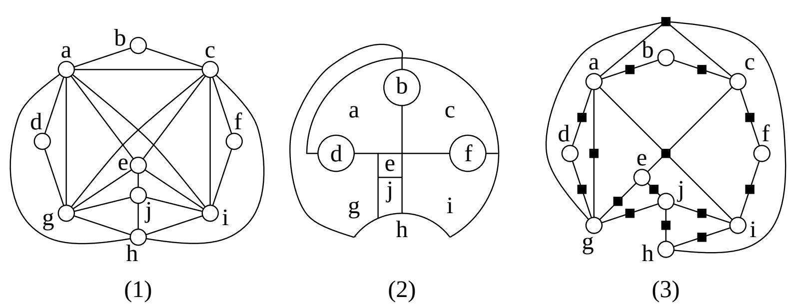

Figure 1.3: A symmetric map graph, a map, and a witness.

1.3 Summary of Results

In the preliminary version of this paper [4] we gave an NP-characterization of map graphs, and we also sketched a polynomial-time recognition algorithm for 4-map graphs. In this paper we prove the first result, but for reasons of brevity we present a simpler variant of the second result. Specifically, we present a polynomial-time recognition algorithm for *hole-free* 4-map graphs. We believe that our argument for the case with holes is correct, however it is a very long case analysis which would double the length of the present paper.

We left general map graph recognition as an open problem; Thorup [15] has recently presented a polynomial-time algorithm for recognizing map graphs. Thorup's result does not necessarily imply ours, since even if we are given a map realizing a map graph, it is not clear that it helps us to find a map with the additional restrictions we want (a hole-free 4-map). Also, Thorup's algorithm is complex and the exponent of its time bound polynomial is about 120, while our algorithm is more understandable and its time bound is $O(n^3)$.

There is an obvious naive approach for all these recognition problems: for each maximal clique in the given graph, assert a point in the map where the nations of the clique should meet. This approach fails because there are other maps which realize a clique, as we will see in Figure 2.1.

In Section 2 we characterize map graphs as the "half-squares" of planar bipartite graphs (Theorem 2.2); this implies that map graph recognition is in NP (Corollary 2.4). Using this, we list all possible clique maps (Theorem 2.7), and we show that a map graph has a linear number of maximal cliques (Theorem 2.8).

In Section 3 we give a high-level presentation of our algorithm for recognizing hole-free 4-map graphs; the algorithm produces such a map, if one exists. In Sections 4 through 7 we present the structural results needed to prove the correctness of the algorithm; these sections are the technical core of our paper. We give a time analysis in Section 8, and concluding remarks in Section 9.

2 Map Graph Fundamentals

2.1 A Characterization

We characterize map graphs in terms of planar bipartite graphs. For a graph H and a subset A of its vertices, let $H[A]$ denote the subgraph of H induced by A. Let H^2 denote the square of H, that is the simple graph with the same vertex set, where vertices are adjacent whenever they are connected by a two-edge path in H. We represent a bipartite graph as $H = (A, B; E)$, where the vertices partition into the independent sets A and B, and E is the set of edges.

Definition 2.1 *Suppose $H = (A, B; E)$ is a bipartite graph. Then $H^2[A]$ is the* half-square *of H. That is, the half-square has vertex set A, where two vertices are adjacent exactly when they have a common neighbor in B.*

Theorem 2.2 *A graph G is a map graph if and only if it is the half-square of some planar bipartite graph H.*

Proof: For the "only if" part, suppose G is a map graph. Let $\mathcal{R}$ be the set of nations in a map of G; for convenience we identify the n vertices of G with the corresponding nations in $\mathcal{R}$.

Consider a single nation R. Clearly at most $n-1$ boundary points will account for all the adjacencies of R with other nations, and so a finite collection $\mathcal{P}$ of boundary points witnesses all the adjacencies among the nations in $\mathcal{R}$.

In each nation R we choose a representative interior point, and connect it with edges through the interior of R to the points of $\mathcal{P}$ bounding R. In this way we construct a planar embedding of the bipartite graph $H = (\mathcal{R}, \mathcal{P}; E')$, such that any two nations R_1 and R_2 overlap if and only if they have distance two in H. In other words, G is the half-square $H^2[\mathcal{R}]$.

For the "if" part, given a bipartite planar graph $H = (\mathcal{R}, \mathcal{P}; E')$, we embed it in the plane. By drawing a sufficiently thin star-shaped nation around each $R \in \mathcal{R}$ and its edges in E', we obtain a map for $H^2[\mathcal{R}]$. $\square$

When graphs G and H are related as above, H acts as a proof that G is a map graph. We call H a *witness* for G, and we call the vertices in $\mathcal{P}$ the *points* of the witness; such points are displayed as squares in the example Figure 1.3(3). The above argument shows that H has at most quadratic size, but we can do better.

Lemma 2.3 *If G is a map graph with n vertices, then it has a witness H with $O(n)$ vertices and edges.*

Proof: Construct H as above. A point $p \in \mathcal{P}$ is *redundant* if all pairs of its neighbors are also connected through other points of $\mathcal{P}$. Deleting a redundant point does not change the half-square; we repeat this until H has no redundant points.

Consider a drawing of H. For each $p \in \mathcal{P}$, we choose a pair of nations R_1 and R_2 connected only by p. Remove each p and its arcs, and replace them by a single arc from R_1 to R_2. In this way, we draw a simple planar subgraph H' of G with edge set $\mathcal{P}$ and vertex set $\mathcal{R}$. Hence $|\mathcal{P}| \leq 3n - 6$, and H has less than $4n$ vertices.

Since H is simple and bipartite, by Euler's formula it has at most $2(n + |\mathcal{P}|) - 4$ edges, which is less than $8n$. $\square$

In particular, a map graph has a witness which may be checked in linear time [10], and so we have:

Corollary 2.4 *The recognition problem for map graphs is in NP.*

Let $\alpha(G)$ denote the *arboricity* of a graph G, the minimum number of forests whose union is G. The next result is useful for the time analysis in Section 8.

Corollary 2.5 *A k-map graph with n vertices has $O(kn)$ edges and arboricity $O(k)$.*

Proof: By Lemma 2.3, the map graph has a witness H with less than $8n$ edges. So, some nation R has degree at most 7 in H, and consequently R has degree less than $7k$ in the map graph. Now we delete R, and prove our edge bound by induction on n.

Since $\alpha(G) = \max_U \lceil \frac{|E(G[U])|}{|U|-1} \rceil$ where U ranges over all subsets of $V(G)$ containing at least two vertices [13], and each $G[U]$ is again a k-map graph, the edge bound implies the arboricity bound. $\square$

For the simplicity of our figures, we prefer to draw maps rather than planar bipartite witness graphs. The arguments in Theorem 2.2 show efficient transformations back and forth between them; so nothing will be lost by using maps. We conclude with some further simple consequences of Theorem 2.2:

- In the witness graph, a point of degree three may be replaced by three points of degree two. Consequently, 3-map graphs are 2-map graphs (planar graphs).
- For $k \geq 3$, k-map graphs are those with witnesses H such that every point has degree at most k.
- Hole-free map graphs are those with witnesses H such that every face has exactly four sides. Since a six-sided face may always be partitioned into three four-sided faces (by adding a redundant point), we may also allow six-sided faces.
- If G has no clique of size four, then it is a map graph if and only if it is a planar graph.
- A map graph may contain cliques of arbitrary size.
- From the previous two remarks, it is clear that the "map graph" property is not monotone, and hence cannot be characterized by forbidden subgraphs or minors.

Regarding the last point, we can also show a stronger example:

Claim 2.6 *There is a hole-free 4-map graph G and edge e such that $G - e$ has no map.*

Proof: Let G be the graph realized by the hole-free 4-map in Figure 1.2(1), and let H be a witness of G. We may assume that H has no point of degree 1. Let H' be the bipartite graph in Figure 1.2(2), and let H'' be the graph obtained from H by deleting points p and q and their incident edges. Since each $v \in \{a, \ldots, j\}$ has exactly two neighbors in G, the degree of v in H is either 1 or 2; we may assume the former case because in the latter case, the two points adjacent to nation v in H can be identified. By this assumption, H'' is an induced subgraph of H. In turn, by the existence of edges $\{1, 6\}$ and $\{4, 6\}$ in G and the planarity of H, H' must be a (not necessarily induced) subgraph of H. Now, by Figure 1.2(2) and planarity, point q must also be adjacent to nations 3 and 7 in H. Our argument so far did not depend on edge $e = \{6, 7\}$ in G, but now this edge has been forced by considering other edges. So in other words, $G - e$ is not a map graph. □

2.2 Cliques in Map Graphs

Suppose G is the clique K_n, then it may be realized in the following ways, corresponding to the four parts of Figure 2.1:

(1) The n nations share a single boundary point. We call this the *pizza*.

(2) Some $n - 1$ nations share a single boundary point, and the one remaining nation is arbitrarily connected to them at other points. We call this the *pizza-with-crust*.

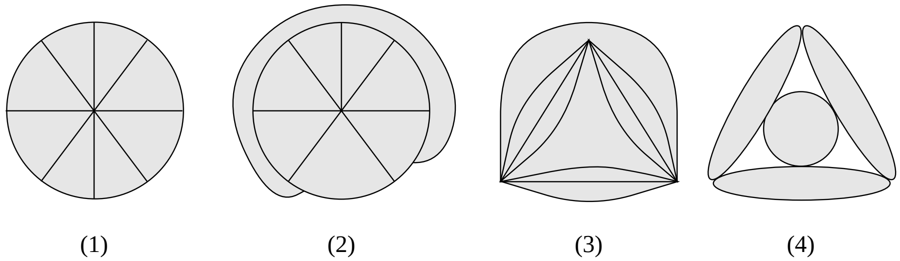

Figure 2.1: Cliques in map graphs.

(3) If $n \geq 6$, there may be three points supporting all adjacencies in the clique, with at most $n-2$ nations at any one point. In particular, there are at most two nations adjacent to all three of the points. We call this the *hamantasch*.

(4) A clique with all boundary points of degree two; that is, an ordinary planar clique. Since the planar K_2 (edge) and K_3 (triangle) are a pizza and pizza-with-crust respectively, the only new clique to list here is the planar K_4, which we call the *rice-ball*.

Theorem 2.7 *A map graph clique must be one of the above four types.*

Proof: By Theorem 2.2, we have a bipartite planar witness graph $H = (\mathcal{R}, \mathcal{P}; E')$ such that the half-square $H^2[\mathcal{R}]$ is the clique K_n, where $n = |\mathcal{R}|$. Let d be the maximum degree of all points $p \in \mathcal{P}$.

If $n = d$, we have a pizza. If $n = d + 1$, we have a pizza-with-crust. So we may assume $n \geq d + 2$. If $d \leq 3$, then the map graph is planar; since K_5 is not planar, this forces $n = 4$ and $d = 2$, the rice-ball. We now assume $d \geq 4$.

Pick point p_1 of maximum degree d, and nations x and y not adjacent to p_1. Consider the set $\mathcal{P}'$ of all points connecting x or y to the nations around p_1. We claim that there is a point $p_2 \in \mathcal{P}'$ connecting x, y, and at least two nations adjacent to p_1; otherwise, by drawing arcs through the points of $\mathcal{P}'$, we could get a planar $K_{d,2}$ with the d nations on a common face, which is impossible. Since p_1 has maximum degree, there are also two nations adjacent to p_1 but not p_2. In summary, the following three disjoint sets each contain at least two nations:

$$\begin{aligned} \mathcal{R}_1 &= \{R \in \mathcal{R} \mid R \text{ is adjacent to } p_2 \text{ but not } p_1\} \\ \mathcal{R}_2 &= \{R \in \mathcal{R} \mid R \text{ is adjacent to } p_1 \text{ but not } p_2\} \\ \mathcal{R}_3 &= \{R \in \mathcal{R} \mid R \text{ is adjacent to both } p_1 \text{ and } p_2\} \end{aligned}$$

We will choose six distinct nations $R_1, R_2 \in \mathcal{R}_1$, $R_3, R_4 \in \mathcal{R}_2$, and $R_5, R_6 \in \mathcal{R}_3$; no matter how we choose, the graph H will contain the induced subgraph in Figure 2.2(1), with the cycle $C = p_1 R_5 p_2 R_6$. The graph $H' = H[\{p_1, p_2\} \cup \mathcal{R}_3]$ is a complete bipartite graph, with a planar embedding inherited from H. Each face of H' is a 4-cycle; furthermore all nations in $\mathcal{R}_1 \cup \mathcal{R}_2$ must lie inside one face, in order to be connected by other points. So, we choose $R_5, R_6 \in \mathcal{R}_3$ on this face. Then this face is bounded by C; we have an embedding with $\mathcal{R}_1 \cup \mathcal{R}_2$ inside C, and $\mathcal{R}_3 - \{R_5, R_6\}$ outside C. By an appropriate choice of nations $R_1, R_2 \in \mathcal{R}_1$ and $R_3, R_4 \in \mathcal{R}_2$, we arrive at Figure 2.2(2), the embedding of p_1,

p_2, and all their edges to adjacent nations. In this figure, the three occurrences of "..." locate any other nations in $\mathcal{R}_1 \cup \mathcal{R}_2 \cup \mathcal{R}_3$.

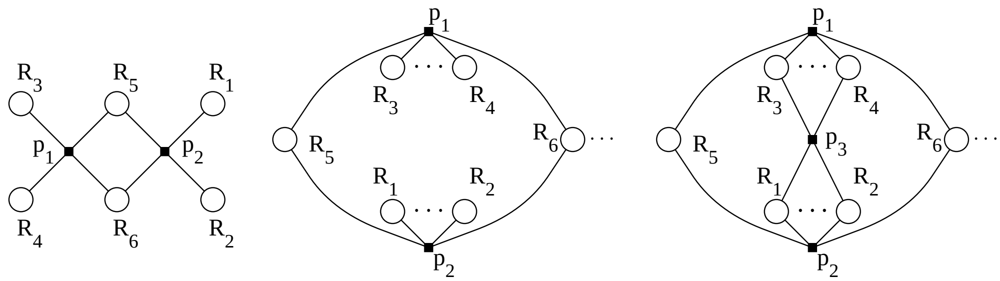

Figure 2.2: A subgraph of H, and its embedding.

There must exist a third point p_3 inside C connecting R_1 and R_4. These edges now separate $\mathcal{R}_1 - \{R_1\}$ from $\mathcal{R}_2 - \{R_4\}$, so all these nations are adjacent to p_3 as well, yielding Figure 2.2(3). This figure is not necessarily an induced subgraph, since the edges $\{R_5, p_3\}$ and $\{R_6, p_3\}$ may occur in H. But by the maximality of the degree of p_1, if exactly $i \in \{1, 2\}$ of these edges exist, then there exist i other nations $R_7 \in \mathcal{R}_3$, necessarily outside C. So, no matter whether these edges exist or not, the points p_1, p_2, and p_3 support a hamantasch on $\mathcal{R}_1 \cup \mathcal{R}_2 \cup \mathcal{R}_3$. Hence, we are done if $\mathcal{R} = \mathcal{R}_1 \cup \mathcal{R}_2 \cup \mathcal{R}_3$.

For contradiction, suppose $\mathcal{R}$ contains some nation R not adjacent to p_1 or p_2. We need to place R, and some new points and edges, in Figure 2.2(3) so that R has neighbor points connected to the other nations. However by planarity of H, if $\{R_5, p_3\}$ or $\{R_6, p_3\}$ is an edge in H, then R cannot be placed so that both R_1 and R_7 have neighbor points connected to R. Similarly, if neither $\{R_5, p_3\}$ nor $\{R_6, p_3\}$ is an edge in H, then R cannot be placed so that all of R_1, R_5, and R_6 have neighbor points connected to R. □

By a careful analysis of each kind of clique, we can now show:

Theorem 2.8 *A map graph G with n vertices has at most $27n$ maximal cliques.*

Proof: We may assume that G is connected. As in Theorem 2.3, we choose a planar witness $H = (\mathcal{R}, \mathcal{P}; E')$ for G where $\mathcal{R}$ is the set of nations, $\mathcal{P}$ is the set of at most $3n - 6$ points, and E' is the set of edges.

Fix a plane embedding of H. If vertices u_1, u_2, v_1, v_2 appear in that cyclic order as distinct neighbors of some vertex w in H, then we say that the pairs $\{u_1, v_1\}$ and $\{u_2, v_2\}$ *cross* at w.

Each point can contribute to at most one maximal pizza, and so there are at most $3n - 6$ maximal pizzas in G. Note that each MC_2 is a pizza.

Next, let $C_1, \ldots, C_\ell$ be the maximal cliques in G that are either non-pizza MC_3's or hamantaschen. For each hamantasch C_i, we may choose three points $p_i, q_i, r_i \in \mathcal{P}$ and three nations $a_i, b_i, c_i \in C_i$ such that $T_i = p_i a_i q_i b_i r_i c_i$ is an induced cycle in H and C_i consists of all nations adjacent to at least two of the points p_i, q_i, r_i in H. For each pair $\{a, b\}$ of vertices in G such that some non-pizza MC_3 contains both a and b, let $s_{a,b}$ be a point in $\mathcal{P}$ that is adjacent to both a and b in H. For each non-pizza MC_3 C_i, let $C_i = \{a_i, b_i, c_i\}$, $p_i = s_{a_i,c_i}$, $q_i = s_{a_i,b_i}$, $r_i = s_{b_i,c_i}$, and T_i be the induced cycle $p_i a_i q_i b_i r_i c_i$ in H. In either case, we define $f(C_i) = \{p_i, q_i, r_i\}$ and note that C_i consists of all nations adjacent to at least two points of $f(C_i)$. This implies $f(C_i) \neq f(C_j)$ for distinct C_i and C_j, because otherwise $C_i \cup C_j$ would be a larger clique.

Define H' as the simple graph with vertex set $\mathcal{P}$ and edge set $\{\{p,q\} \mid \{p,q\} \subset f(C_i)$ for some $i\}$. We claim that H' is planar. To see this, we embed H' in H by drawing the edge $\{p_i, q_i\}$ of H' through their neighbor a_i in H, and similarly for the other two edges $\{q_i, r_i\}$ and $\{p_i, r_i\}$. Towards a contradiction, assume that two edges of H' cross in the embedding. Then, by cycle symmetries, we may assume that for some distinct C_i and C_j, pairs $\{p_i, q_i\}$ and $\{p_j, q_j\}$ cross at nation $a_i = a_j$ (call it a) in H. Since the cycles T_i and T_j cross at a in H, they must cross again, sharing either another nation or a point.

Case 1: T_i and T_j share another nation but no point. By symmetry, it suffices to consider the case $b_i = b_j = b$. If C_i and C_j were both non-pizza MC_3's, then we would have $q_i = q_j = s_{a,b}$, contradicting the crossing. So at least one of C_i and C_j is a hamantasch, we suppose C_i. Then C_i has another nation a' also adjacent to p_i and q_i or else C_i would be a pizza-with-crust. Because of T_j, it must be $a' = c_j$. Now since q_i is adjacent to a, b, c_j, in order for C_j to be a non-pizza, C_j must also be a hamantasch. Then C_j has another nation a'' adjacent to p_j and q_j, but planarity of H makes this impossible.

Case 2: T_i and T_j share a point. Since T_i and T_j are induced, a is adjacent to neither r_i nor r_j, so the only possible shared point is $r_i = r_j$. In turn, $C = \{a, b_i, c_i, b_j, c_j\}$ is a clique of G. So, neither C_i nor C_j is an MC_3, and both are hamantaschen. Now as in the previous case, we find it is impossible to add a nation $a' \notin C$ between p_i and q_i and a nation $a'' \notin C$ between p_j and q_j. By this, both C_i and C_j are pizza-with-crusts, a contradiction.

By the above case-analysis, the claim holds, and H' is a planar graph with at most $3n-6$ vertices and at least ℓ distinct triangles. An easy exercise shows that any simple planar graph with h vertices has at most $3h$ triangles. So, $\ell \leq 9n - 18$.

There are at most n rice-balls, since they all have different center nations.

It remains to bound the number of maximal pizza-with-crusts of size 4 or more. Fix a point p in H, let V_p denote the set of nations adjacent to p in H, and let $\mathcal{C}_p = \{C_1, \ldots, C_{\ell_p}\}$ be the maximal pizza-with-crusts with center p and size 4 or more. We claim that $\ell_p = |\mathcal{C}_p| \leq 2(2|V_p| - 3)$. This claim implies that G has at most $\sum_p (4|V_p| - 6)$ maximal pizza-with-crusts of size 4 or more. This sum equals $4|E'| - 6|\mathcal{P}|$; since $|E'| \leq 2(n + |\mathcal{P}|) - 4$ and $|\mathcal{P}| \leq 3n - 6$, the sum is less than $14n$.

Now we prove the claim. The embedding gives a cyclic clockwise order on the nations of V_p around p; this order defines "consecutive" nations and "intervals" of nations around p. For nations $u, v \in V_p$, let $[u, v]$ denote the circular interval of nations starting at u, proceeding clockwise around p, and ending at v. For each clique C_i in $\mathcal{C}_p$, let b_i be the crust of C. Since C_i is not a pizza, we can choose distinct nations $a_i, c_i \in C_i - \{b_i\}$ and distinct points $q_i, r_i \neq p$ satisfying the following three conditions:

- $T_i = p a_i q_i b_i r_i c_i$ is a simple cycle in H.
- If a nation of $C_i - \{a_i, b_i, c_i\}$ is adjacent to q_i or r_i, then it lies outside T_i, otherwise it lies inside.
- All nations of C_i lying inside the cycle T_i lie in the interval $[a_i, c_i]$.

Denote the unordered pair $\{a_i, c_i\}$ by $g(C_i)$, and $\{q_i, r_i\}$ by $h(C_i)$. By considering such 6-cycles T_i and the planarity of H, there are no cliques $C_i, C_j \in \mathcal{C}_p$ such that $g(C_i)$ and

$g(C_j)$ cross at p; consequently the graph $G_p = (V_p, \{g(C_i) \mid C_i \in \mathcal{C}_p\})$ is simple outerplanar, where we use the same cyclic order on V_p for the outerplanar embedding. Since G_p is simple outerplanar, it can have at most $2|V_p|-3$ edges. Thus, to prove $|\mathcal{C}_p| \leq 2(2|V_p|-3)$, it suffices to prove that each edge of G_p equals $g(C_i)$ for at most two $C_i \in \mathcal{C}_p$.

For contradiction, assume that there exist three distinct cliques $C_i, C_j, C_k \in \mathcal{C}_p$ with $g(C_i) = g(C_j) = g(C_k)$. Say that two of these cliques are *nested* if the crust of one is inside the cycle of the other. We consider two cases.

Case I: There are two non-nested cliques. We may suppose that they are C_i and C_j. By planarity, the interiors of T_i and T_j are disjoint, with $a_i = b_j$ and $a_j = b_i$. Moreover, no matter whether $h(C_i) \cap h(C_j) = \emptyset$ or not, no nation of $V_p - g(C_i)$ is adjacent to both p and at least one point of $h(C_i) \cup h(C_j)$ in H. Thus, the set of nations lying inside T_i and the set of nations lying inside T_j form a partition of $V_p - g(C_i)$. By this and the maximality of cliques in $\mathcal{C}_p$, the crust of C_k must lie inside T_i or T_j; by symmetry we suppose it lies inside T_i. On the other hand, since C_i is a maximal clique of size 4 or more, there exists a nation $x_i \in C_i - g(C_i)$ lying inside T_i. Since x_i cannot be adjacent to q_i or r_i, there is another point s_i lying inside T_i that connects x_i with b_i. So, we have a path $P_i = px_is_ib_i$ sharing only its endpoints with T_i, and bisecting the interior of T_i. Now the crust b_k must lie inside T_i, to one side or the other of P_i. To achieve $g(C_k) = g(C_i)$, b_k must be adjacent to s_i in H. But then we would see that $s_i \in h(C_k)$, and so either a_k or c_k was chosen incorrectly.

Case II: All three cliques nest. Again by planarity, the cycles cannot cross. So, their interiors nest in some order; we may assume that b_k lies inside T_j, and b_j lies inside T_i. We have $a_i = a_j = a_k$ and $c_i = c_j = c_k$. Since C_j is a maximal clique, it contains some nation x_j not adjacent to b_i in G. By planarity, x_j must lie inside T_j, and there is some point s_j connecting b_j with x_j in H. We cannot have $s_j \in h(C_j)$, by the choice of $g(C_j)$; so again we have a path $P_j = px_js_jb_j$, sharing only its endpoints with T_j and bisecting its interior. Now the crust b_k must lie inside T_j, to one side or the other of P_j; the rest of the argument proceeds as in the last case. □

3 Recognizing Hole-Free 4-Map Graphs

The rest of this paper is devoted to an $O(n^3)$-time algorithm for deciding whether a given graph is a hole-free 4-map graph. It follows from Theorem 2.7 that 4-map graphs have no 7-cliques, that all 6-cliques are hamantaschen, and that all 5-cliques are pizza-with-crusts. Also, as observed in Section 1.2, all hole-free map graphs are 2-connected. Unfortunately, these simple observations do not lead to a polynomial-time algorithm. Indeed, our algorithm is very sophisticated and too long to be included in a single section. This section only gives a high-level sketch of the algorithm. The correctness and implementation details are given in subsequent sections.

3.1 Definitions

Before presenting the algorithm, we define some vocabulary for the remainder of the paper. A *marked graph* is a simple graph in which each edge is either marked or not marked. Throughout the remainder of this paper, G is a marked graph with vertex set $V(G) = V$ and edge set $E(G) = E$. For a vertex v in G, $N_G(v)$ denotes the set of vertices adjacent to v in G. For a subset U of V, $N_G(U)$ denotes $\cup_{u \in U} N_G(u)$. Let F be a subset of E. $G - U - F$ denotes the marked graph obtained from G by deleting the edges in F and the vertices in U together with the edges incident to them. When U or F is empty, we drop it from the notation $G - U - F$.

Definition 3.1 *Let U be a subset of V. A* layout $\mathcal{L}$ *of $G[U]$ is a 4-map $\mathcal{L}$ of $G[U]$ such that for every marked edge $\{u, v\}$ in G with $u, v \in U$, the boundaries of nations $\mathcal{L}(u)$ and $\mathcal{L}(v)$ share a curve segment (not just one or more isolated points). $\mathcal{L}$ is* well-formed *if for every edge $\{u, v\}$ in $G[U]$, the intersection of $\mathcal{L}(u)$ and $\mathcal{L}(v)$ is either a point or a curve segment (but not both).*

Definition 3.2 *If a layout $\mathcal{L}$ of G covers every point of the sphere, we call it an* atlas *of G.*

Our goal is to design an efficient algorithm to decide whether a given G has an atlas; it will either return an atlas or report failure. Furthermore, the algorithm returns a well-formed atlas whenever possible (see Corollary 4.2).

If G has an atlas, then by Lemma 2.3 it has a witness graph checkable in linear time [10]. So in fact we describe an algorithm that makes the following assumption:

Assumption 1 *G has an atlas.*

If G does not have an atlas, we will discover this when our algorithm either fails, returns an invalid atlas, or takes more time than allowed by our analysis in Section 8. Since we assume that G has an atlas, we will call the vertices in G *nations*; we will use lower-case letters to denote them.

Throughout the rest of this subsection, fix a subset U of V and a layout $\mathcal{L}$ of $G[U]$. A vertex $u \in U$ *touches* a hole $\mathcal{H}$ of $\mathcal{L}$ if $\mathcal{L}(u)$ intersects $\mathcal{H}$; they necessarily intersect on their boundaries. Vertex u *strongly touches* $\mathcal{H}$ if the boundaries of $\mathcal{L}(u)$ and $\mathcal{H}$ share a curve segment. A *2-hole* is a hole strongly touched by exactly two vertices. *Erasing a 2-hole $\mathcal{H}$* in $\mathcal{L}$ is the operation of modifying $\mathcal{L}$ by extending $\mathcal{L}(u)$ to occupy $\mathcal{H}$, where u is one of the vertices strongly touching $\mathcal{H}$. Figure 3.1(1) depicts the operation.

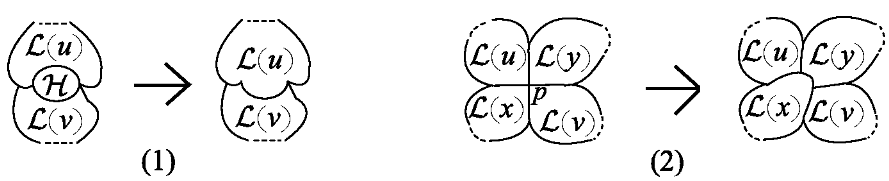

Figure 3.1: Erasing a 2-hole $\mathcal{H}$, and a (u, v)-point p. Dashed curves may intersect.

A *k-point* in $\mathcal{L}$ is a point shared by exactly k nations. Let $u \in U$ and $v \in U$. A *(u, v)-point* in $\mathcal{L}$ is a 4-point p at which $\mathcal{L}(u)$ and $\mathcal{L}(v)$ together with two other nations $\mathcal{L}(x)$ and $\mathcal{L}(y)$ meet cyclically in the order $\mathcal{L}(u)$, $\mathcal{L}(x)$, $\mathcal{L}(v)$, $\mathcal{L}(y)$. *Erasing the (u, v)-point p* in $\mathcal{L}$ is the operation of modifying $\mathcal{L}$ by extending nation $\mathcal{L}(x)$ to occupy a small disc around p touching only $\mathcal{L}(u)$, $\mathcal{L}(v)$, $\mathcal{L}(x)$, and $\mathcal{L}(y)$. Figure 3.1(2) depicts the operation.

A *(u, v)-segment* in $\mathcal{L}$ is a curve segment S shared by the boundaries of $\mathcal{L}(u)$ and $\mathcal{L}(v)$ such that each endpoint of S is a 3- or 4-point. Note that two (u, v)-segments must be disjoint. An edge $\{u, v\}$ of G is *good* in $\mathcal{L}$ if there is either exactly one (u, v)-segment or exactly one (u, v)-point, but not both. An edge that is not good in $\mathcal{L}$ is *bad* in $\mathcal{L}$. Note that $\mathcal{L}$ is well-formed if and only if every edge of $G[U]$ is good in $\mathcal{L}$.

Definition 3.3 *If $\mathcal{M}$ is an atlas of G and W is a subset of V, then $\mathcal{M}|_W$ denotes the layout of $G[W]$ obtained by restricting $\mathcal{M}$ to W. Such a layout is called an* extensible *layout of $G[W]$. $\mathcal{L}$ is* transformable *to another layout $\mathcal{L}'$ of $G[W]$ if whenever $\mathcal{L}$ is extensible, so is $\mathcal{L}'$.*

Note that an extensible layout never has a 4-point on a hole boundary, since filling the hole would create an illegal 5-point. Similarly, if two holes touch in an extensible layout, it must be at a 2-point.

3.2 Making Progress

To find an atlas for G, our algorithm may "make progress" by producing one or more smaller marked graphs, so that finding an atlas for G is reduced to finding an atlas for each of these smaller graphs. Here we define the graph features that our algorithm may identify in order to make progress; subsequent sections show how to make progress for each.

A *k-cut* of G is a subset U of V with $|U| = k$ whose removal disconnects G. G is *k-connected* if it has no i-cut with $i \leq k - 1$. Section 4 shows that the algorithm can always make progress when G is not 4-connected. On the other hand, under the assumption that G is 4-connected, Corollary 4.2 guarantees that G has a well-formed atlas.

Definition 3.4 *A clique consisting of k vertices is called a k-*clique*. A clique C in G is* maximal *if no clique in G properly contains C. A maximal k-clique is called an MC_k.*

If G is 4-connected and has at most 8 vertices, our algorithm will construct a well-formed atlas for G by exhaustive search. On the other hand, under the assumptions that G is 4-connected and has at least 9 vertices, Lemma 4.4 guarantees that G has no 6-clique.

Definition 3.5 *A* correct 4-pizza *is a list $\langle a, b, c, d\rangle$ of four nations in G such that G has a well-formed atlas in which nations a, b, c, d meet at a point cyclically in this order.* Removing *a correct 4-pizza $\langle a, b, c, d\rangle$ from G is the operation of modifying G as follows: Delete the edge $\{a, c\}$ from G and mark the edges $\{b, d\}$, $\{a, b\}$, $\{b, c\}$, $\{c, d\}$, and $\{d, a\}$.*

Under the assumption that G is 4-connected, Lemma 4.5 allows our algorithm to make progress by removing a correct 4-pizza in G, whenever we identify one.

To see a particular type of correct 4-pizza, consider an extensible layout of an MC_5 C in G. Since a 5-clique is not planar, the layout contains at least one 4-point. Inspection shows there is exactly one 4-point in a well-formed layout. This "pizza-with-crust" layout motivates the following definition.

Definition 3.6 *A* correct center *of C is a list $\langle a, b, c, d\rangle$ of four nations in C, such that C has a well-formed extensible layout in which nations a, b, c, d meet at a point in this cyclic order. The unique nation in $C - \{a, b, c, d\}$ is the corresponding* correct crust *of C.*

Fact 3.7 *Let C be an MC_5 in G. Then, every correct center of C is a correct 4-pizza in G.*

Note that C may have multiple correct centers, each from a different extensible layout.

Besides the k-cuts mentioned above, we also consider the more specialized separators introduced below in Definition 3.10. Section 5 will show how the algorithm may make progress as long as G contains one of these.

Definition 3.8 *Edges $\{a, b\}$ and $\{x, y\}$ in G are* crossable *if they are both unmarked and $\{a, b, x, y\}$ is an MC_4 in G. For an edge $\{a, b\}$, let $\mathcal{E}[a, b]$ denote the set of all edges $\{x, y\}$ crossable with $\{a, b\}$.*

Fact 3.9 *If $\{a, b\}$ is an edge and $G - \{a, b\}$ has a triangle $\{c, d, e\}$, then at most one edge of that triangle is in $\mathcal{E}[a, b]$.*

Proof: Two edges would imply two MC_4's, sitting inside the 5-clique $\{a, b, c, d, e\}$. □

Definition 3.10 *We define the following separators in the marked graph G:*

1. *A* separating edge *of G is an edge $\{a, b\}$ such that $G - \{a, b\} - \mathcal{E}[a, b]$ is disconnected.*
2. *An* induced 4-cycle *in G is a set C of four vertices in G such that $G[C]$ is a cycle. A* separating 4-cycle *of G is an induced 4-cycle C in G such that $G - C$ is disconnected.*
3. *A* separating triple *of G is a list $\langle a, b, c\rangle$ of three vertices in G such that $C = \{a, b, c\}$ is a clique in G and $G - C - \mathcal{E}[a, b]$ is disconnected.*
4. *A* separating quadruple *is a list $\langle a, b, c, d\rangle$ of four vertices in G such that (i) $G[\{a, b, c, d\}]$ is a cycle and (ii) $G - \{a, b, c, d\} - \mathcal{E}[a, b]$ is disconnected.*
5. *A* separating triangle *of G is a list $\langle a, b, c\rangle$ of three vertices in G such that (i) $C = \{a, b, c\}$ is a clique in G and (ii) $G' = G - C - (\mathcal{E}[a, b] \cup \mathcal{E}[a, c])$ is disconnected. If in addition, G' has a connected component consisting of a single vertex, then $\langle a, b, c\rangle$ is a* strongly separating triangle *of G.*

3.3 Sketch of the Algorithm

Given a marked graph $G = (V, E)$, our algorithm rejects if G is not 2-connected, and it solves the problem by exhaustive search when $|V| \leq 8$. When $|V| \geq 9$, it searches G for a 2-cut (cf. Lemma 4.1), 3-cut (cf. Lemma 4.3), separating edge (cf. Lemma 5.2), separating 4-cycle (cf. Lemma 5.4), separating triple (cf. Lemma 5.5), separating quadruple (cf. Lemma 5.6), strongly separating triangle (cf. Lemma 5.14), or separating triangle (cf. Lemma 5.15), *in this order*. In each case, as the lemmas show, the algorithm makes progress by either (1) removing a correct 4-pizza or (2) reducing the problem for G to the problems for certain marked graphs smaller than G whose total size is that of G plus a constant.

If none of the above separators exists in G, then G has no 6-clique (cf. Lemma 4.4) and the algorithm searches G for an MC_5 or MC_4, *in this order*. If an MC_5 C is found, it tries to find an extensible layout of C by doing a case-analysis based on the neighborhood of C

in G (cf. Section 6). The absence of the above separators guarantees that only a few cases needed to be analyzed. The case-analysis either yields an extensible layout of C whose center is then removed to make progress, or produces a marked graph G' smaller than G such that finding a well-formed atlas for G can be reduced to finding a well-formed atlas for G'.

If no MC_5 but an MC_4 is found in G, the algorithm scans all MC_4's of G in an arbitrary order. While scanning an MC_4 C, it decides whether C has a rice-ball layout (cf. Lemma 7.1). If C has a rice-ball layout, the algorithm quits the scanning and makes progress by removing a correct 4-pizza obtained from the rice-ball layout of C. On the other hand, if no rice-ball is found after scanning all MC_4's, the algorithm scans all MC_4's of G in an arbitrary order, once again. But this time, while scanning an MC_4 C, it decides whether C has a non-pizza layout, by doing a case-analysis based on the neighborhood of C in G (cf. Section 7.2). The analysis consists of only a few cases due to the absence of the above separators. If C has a non-pizza layout, the algorithm quits the scanning and makes progress by removing a correct 4-pizza obtained from the layout of C. Otherwise, all MC_4's are pizzas; the algorithm finds their centers (cf. Section 7.3), and removes all of them so that G no longer has an MC_4.

If neither MC_5 nor MC_4 is found in G, then this is a base case. As observed in Section 2.1, G must be planar, or else we reject. When G is planar, then by its 4-connectivity it has a unique planar embedding. We claim that G has a well-formed atlas if and only if all its faces are triangles. The "if" direction is obvious because the planar dual of G is an atlas, which is well-formed by the connectivity of G. Conversely, suppose G has a well-formed atlas $\mathcal{M}$. Since $\mathcal{M}$ has no k-point for $k > 3$, all adjacent pairs of nations strongly touch in $\mathcal{M}$, and so the 3-points and boundaries in $\mathcal{M}$ define a 3-regular planar graph G', whose dual is G. So, it suffices for the algorithm to check that G is planar and has a 3-regular dual; if so, it returns the dual as an atlas.

In all the recursive cases, the smaller graphs that we generate have total size at most the size of G plus a constant, and we spend quadratic time on generating them. A simple argument (cf. Section 8) shows that the overall time is cubic.

3.4 Figures

For our arguments of the algorithm's correctness, we need a convenient graphical notation for the possible extensible layouts of $G[U]$, where U is some small subset of $V(G)$. First, as is very natural, we consider two layouts equivalent when they are homeomorphic. But beyond this, we also introduce a convenient graphic notation for partially determined layouts of $G[U]$. In particular, we introduce contractible segments and permutable labels.

Definition 3.11 *A* figure *of $G[U]$ is a list $\mathcal{D} = \langle \mathcal{L}, \mathcal{S}, L_1, \ldots, L_k \rangle$, where $\mathcal{L}$ is a layout of $G[U]$, $\mathcal{S}$ is a set of curve segments in $\mathcal{L}$, and $L_1, \ldots, L_k$ are disjoint lists of vertices in U. We call $\mathcal{L}$ the* layout *in $\mathcal{D}$, call the curve segments in $\mathcal{S}$ the* contractible segments *in $\mathcal{D}$, and call $L_1, \ldots, L_k$ the* permutable lists *in $\mathcal{D}$.*

We say $\mathcal{D}$ displays *a layout $\mathcal{L}'$ of $G[U]$ if $\mathcal{L}'$ can be obtained from $\mathcal{L}$ by:*

1. *contracting parts of some contractible segments to points,*

2. *erasing all resulting 2-holes, and*

3. for each permutable list L_i, selecting a permutation π of L_i and renaming each nation $u \in L_i$ as $\pi(u)$.

We say $\mathcal{D}$ displays *$G[U]$, or $\mathcal{D}$ is a* display *of $G[U]$, if $\mathcal{D}$ displays an extensible layout of $G[U]$. $\mathcal{D}$ is* transformable *to another figure $\mathcal{D}'$ of $G[U]$ if whenever $\mathcal{D}$ displays $G[U]$, so does $\mathcal{D}'$.*

To illustrate a figure $\mathcal{D}$ we draw $\mathcal{L}$, emphasize the contractible segments in bold, and then for each permutable list L_i, we label each nation $u \in L_i$ as u^i. Since we may contract any part of a contractible segment to a point, the endpoints do not matter, and they are not emphasized. The holes are unlabeled, and should be regarded as "optional" if a contraction could reduce it to a 2-hole.

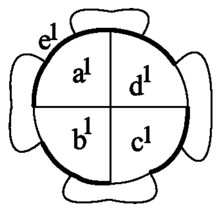

Figure 3.2: A display of MC$_5$ $\{a, b, c, d, e\}$.

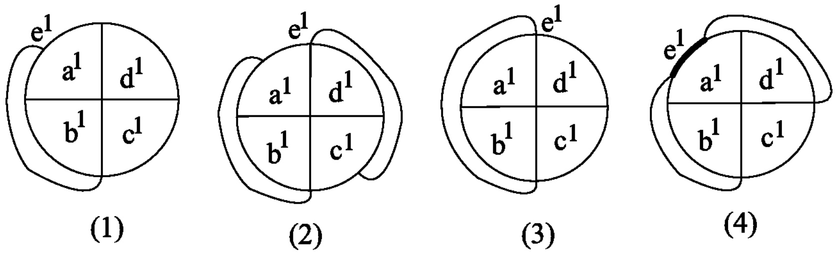

Figure 3.3: Possible displays of MC$_5$ $\{a, b, c, d, e\}$.

For example, when G has a well-formed atlas $\mathcal{M}$ and at least 6 vertices but has no separating triangle, Figure 3.2 displays $G[U]$ for an MC$_5$ $U = \{a, b, c, d, e\}$ of G. To see this, we contract a set of contractible segments in the figure to obtain Figure 3.3(1) through (4). We argue that one of them must display $\mathcal{M}|_U$ as follows. First, U is a pizza-with-crust in $\mathcal{M}$. Suppose that the four non-crust nations a^1, b^1, c^1, d^1 meet at a 4-point p in $\mathcal{M}$ in this order. Let $q_{a,b}$ be the endpoint of the (a^1, b^1)-segment in $\mathcal{M}$ other than p. Define $q_{b,c}$, $q_{c,d}$, and $q_{d,a}$ similarly. Let k be the number of points among $q_{a,b}$, $q_{b,c}$, $q_{c,d}$, $q_{d,a}$ that are touched by the crust e^1 of U and another nation of $V - U$ in $\mathcal{M}$. Since $\mathcal{M}$ is well-formed, $k \leq 2$. On the other hand, since $G \neq G[U]$ and has no separating triangle, $k \geq 1$. If $k = 1$, then Figure 3.3(1) displays $\mathcal{M}|_U$. If $k = 2$, then Figure 3.3(2), (3) or (4) displays $\mathcal{M}|_U$. We note that Figure 3.2 has a unique permutable set, namely, U itself.

4 Establishing Connectivity

Our goal here is to reduce the algorithmic problem to the case when G is 4-connected. Since G is already 2-connected, we first show how to reduce to the 3-connected case.

Lemma 4.1 *Let $\mathcal{M}$ be an atlas of G. Let u and v be two distinct vertices of G. Then, the following statements hold:*

1. *$G - \{u, v\}$ is disconnected if and only if there are at least two (u, v)-segments in $\mathcal{M}$.*
2. *Suppose that $G-\{u, v\}$ is disconnected and its connected components are $G_1, \ldots, G_k$. Then for each $i \in \{1, \ldots, k\}$, the marked graph G'_i obtained from $G[V(G_i) \cup \{u, v\}]$ by marking edge $\{u, v\}$ has an atlas. Moreover, given an atlas $\mathcal{M}_i$ for each G'_i, we can easily construct an atlas of G.*

Proof: We first prove Statement 1. If $\{u, v\} \notin E$, then nations u and v are disjoint disc homeomorphs in $\mathcal{M}$, and hence removing them from $\mathcal{M}$ leaves exactly one connected region. So we know $\{u, v\} \in E$. Let k be the number of (u, v)-segments in $\mathcal{M}$. Consider the following three cases.

Case 1: $k = 0$. We erase the (u, v)-points in $\mathcal{M}$. Then, $\mathcal{M}$ becomes an atlas of $G - \{\{u, v\}\}$ and nations u and v are disjoint disc homeomorphs in it. So, removing u and v from $\mathcal{M}$ leaves exactly one connected region.

Case 2: $k = 1$. We erase the (u, v)-points in $\mathcal{M}$. $\mathcal{M}$ remains an atlas of G. Moreover, edge $\{u, v\}$ becomes good in $\mathcal{M}$. Since the union of nations u and v is a disc homeomorph in $\mathcal{M}$, removing them from $\mathcal{M}$ leaves exactly one connected region.

Case 3: $k \geq 2$. We erase the (u, v)-points in $\mathcal{M}$. $\mathcal{M}$ remains an atlas of G. Moreover, there are exactly k disjoint holes in $\mathcal{M}|_{\{u,v\}}$. So, removing nations u and v from $\mathcal{M}$ leaves exactly k connected regions. Each of these regions forms a connected component of $G - \{u, v\}$. This completes the proof of Statement 1.

We next prove Statement 2. For each i, let $U_i = V(G_i)$. Each hole in $\mathcal{M}|_{U_i \cup \{u,v\}}$ is a 2-hole and is touched only by u and v, and hence erasing the holes in $\mathcal{M}|_{U_i \cup \{u,v\}}$ yields an atlas of G'_i. On the other hand, given an atlas $\mathcal{M}_i$ of each G'_i, we erase any (u, v)-points in $\mathcal{M}_i$. $\mathcal{M}_i$ remains an atlas of G'_i, because edge $\{u, v\}$ is marked in G'_i and so there exists a (u, v)-segment in $\mathcal{M}_i$. Since $G'_i - \{u, v\} = G_i$ is connected, Statement 1 implies there is exactly one (u, v)-segment. Thus removing nations u and v from $\mathcal{M}_i$ leaves exactly one connected region, and the closure $\mathcal{L}_i$ of this region is a disc homeomorph. The boundary of $\mathcal{R}_i$ can be divided into two curve segments $S_{i,u}$ and $S_{i,v}$ such that $S_{i,u}$ (respectively, $S_{i,v}$) is a portion of the boundary of nation u (respectively, v) in $\mathcal{M}$. Now, we can obtain an atlas of G as follows. First, put $\mathcal{R}_1, \ldots, \mathcal{R}_k$ on the sphere in such a way that no two of them touch and each $S_{i,u}$ appears on the upper half of the sphere while each $S_{i,v}$ appears on the lower half. Second, draw nation u (respectively, v) to occupy the area of the upper (respectively, lower) half of the sphere that is occupied by no $\mathcal{R}_i$. This gives an atlas of G. □

Corollary 4.2 *G is 3-connected if and only if G has a well-formed atlas.*

Proof: By Lemma 4.1, the "if" part is obvious. For the other direction, suppose G is 3-connected. Let $\mathcal{M}$ be an atlas of G. If no edge of G is bad in $\mathcal{M}$, then $\mathcal{M}$ is well-formed and we are done. So, suppose that some edge $\{u, v\}$ is bad in $\mathcal{M}$. Since G is 3-connected, there is at most one (u, v)-segment in $\mathcal{M}$. If there is no (u, v)-segment in $\mathcal{M}$, we erase all but one (u, v)-points in $\mathcal{M}$; otherwise, we erase all the (u, v)-points in $\mathcal{M}$. In both cases,

$\mathcal{M}$ remains an atlas of G and edge $\{u, v\}$ becomes good in $\mathcal{M}$ while no good edge becomes bad in $\mathcal{M}$. Consequently, we can make all bad edges good in $\mathcal{M}$. $\square$

Using Statement 2 of Lemma 4.1, our algorithm may reduce the given graph to its 3-connected components, and so we assume:

Assumption 2 *G is 3-connected, and it has a well-formed atlas denoted by $\mathcal{M}$.*

We say that two nations u and v *strongly touch* in $\mathcal{M}$ if there is a (u, v)-segment in $\mathcal{M}$; they *weakly touch* in $\mathcal{M}$ if there is a (u, v)-point in $\mathcal{M}$. To simplify the sequel, we also suppose that small graphs are handled by exhaustive methods. We assume:

Assumption 3 *G has $|V| \geq 9$ vertices.*

Lemma 4.3 *Let $C = \{a, b, c\}$ be a set of three distinct vertices in G. Then, the following statements hold:*

1. *When C is not a clique in G, $G - C$ is connected.*
2. *When C is a clique in G, $G - C$ is disconnected if and only if (i) the nations in C do not meet at a point in $\mathcal{M}$ and (ii) each pair of nations in C strongly touch in $\mathcal{M}$.*
3. *Suppose that $G - C$ is disconnected. Then, (i) $G - C$ has exactly two connected components G_1 and G_2, and (ii) both G'_1 and G'_2 have a well-formed atlas, where G'_1 (respectively, G'_2) is the marked graph obtained from $G[V(G_1) \cup C]$ (respectively, $G[V(G_2) \cup C]$) by marking the edges in $E(G[C])$. Moreover, given a well-formed atlas of G'_1 and one of G'_2, we can easily construct one of G.*

Proof: To prove Statement 1, suppose that C is not a clique. For each edge $\{u, v\} \in E(G[C])$, if nations u and v weakly touch in $\mathcal{M}$, then we erase the (u, v)-point in $\mathcal{M}$. Then, $\mathcal{M}|_{V-C}$ is a layout of $G - C$ and the holes in $\mathcal{M}|_{V-C}$ are disjoint disc homeomorphs. So, $G - C$ must be connected.

To prove Statement 2, suppose that C is a clique. The "if" part is clear. To prove the "only if" part, suppose that (i) or (ii) in Statement 2 does not hold. In case (i) is false, a, b and c meet at a point in $\mathcal{M}$, and the well-formedness of $\mathcal{M}$ ensures that their union is a disc homeomorph, and so $G - C$ is connected. Otherwise, suppose (i) is true and (ii) is false. For each edge $\{u, v\} \in E(G[C])$, if nations u and v weakly touch in $\mathcal{M}$, then we erase the (u, v)-point to get atlas $\mathcal{M}'$. Then $\mathcal{M}'|_{V-C}$ is a layout of $G - C$ and the holes in $\mathcal{M}'|_{V-C}$ are disjoint disc homeomorphs. So, $G - C$ is connected.

Next, we prove Statement 3. Since $G - C$ is disconnected, (i) and (ii) in Statement 2 hold. By this, there are exactly two holes $\mathcal{H}_1$ and $\mathcal{H}_2$ in $\mathcal{M}|_C$ and they are disjoint. For $i \in \{1, 2\}$, let U_i be the set of nations that occupy $\mathcal{H}_i$ in atlas $\mathcal{M}$. Each $G[U_i]$ is a connected component of G. Let G'_i be the marked graph obtained from $G[U_i \cup C]$ by marking the edges in $E(G[C])$. There is a unique hole in $\mathcal{M}|_{U_1 \cup C}$ and it is (strongly) touched only by the nations of C. So, modifying $\mathcal{M}|_{U_1 \cup C}$ by extending nation a to occupy its unique hole yields a well-formed atlas of G'_1. Similarly, we can obtain a well-formed atlas of G'_2.

On the other hand, suppose that we are given a well-formed atlas $\mathcal{M}_1$ of G'_1 and one $\mathcal{M}_2$ of G'_2. Let $i \in \{1, 2\}$. Since the edges in $E(G[C])$ are marked in G'_i, each pair of nations of C strongly touch in $\mathcal{M}_i$. Note that $G'_i - C$ is connected. Then by Statement 2 and the well-formedness of $\mathcal{M}_i$, the nations in C meet at a 3-point p_i in $\mathcal{M}_i$. Let D_i be a disc that is centered at p_i and touches no nation of U_i in atlas $\mathcal{M}_i$. To obtain a well-formed atlas of G, we remove each D_i from $\mathcal{M}_i$ to obtain a connected region $\mathcal{R}_i$, and then glue $\mathcal{R}_1$ and $\mathcal{R}_2$ together by identifying nations a, b, c in $\mathcal{R}_1$ with those in $\mathcal{R}_2$, respectively. □

Statement 3 above gives an effective algorithm to reduce the given graph to its 4-connected components. So we now assume:

Assumption 4 *Graph G is 4-connected.*

Lemma 4.4 *G has no 6-clique.*

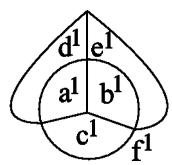

Figure 4.1: A display of MC_6 $\{a, \ldots, f\}$.

Proof: First observe that no hole-free 4-map graph has a 7-clique, by Theorem 2.7. Assume, on the contrary, that G has an MC_6 C. Then it must be a hamantasch, and by Assumption 4, Figure 4.1 displays $\mathcal{M}|_C$. Thus, $V = C$, contradicting Assumption 3. □

Lemma 4.5 *Let G' be the marked graph obtained from G by removing a correct 4-pizza $\langle a, b, c, d\rangle$. Then, G' has a well-formed atlas. Moreover, given a well-formed atlas of G', we can easily construct one of G.*

Proof: Let $\mathcal{M}$ be a well-formed atlas of G in which nations a, b, c, d meet at a point p in this order. After erasing the (a, c)-point p in $\mathcal{M}$, we obtain a well-formed atlas of G' in which nations a, b, and d meet at a 3-point and nations b, c, and d meet at a 3-point. Thus, by Lemma 4.3, both $G' - \{a, b, d\}$ and $G' - \{b, c, d\}$ are connected.

Let $\mathcal{M}'$ be a well-formed atlas of G'. Since $G' - \{a, b, d\}$ is connected and the edges $\{a, b\}$, $\{a, d\}$, and $\{b, d\}$ are marked in G', nations a, b, and d must meet at a 3-point p_1 in $\mathcal{M}'$ according to Lemma 4.3. Similarly, nations b, c, and d must meet at a 3-point p_2 in $\mathcal{M}'$. Thus, the boundaries of b and d in $\mathcal{M}'$ share a curve segment S with endpoints p_1 and p_2. We modify $\mathcal{M}'$ by contracting S to a point, obtaining a well-formed atlas of G. □

In some of our reductions we will discover that G has a well-formed map with several 4-pizzas. In those situations we may remove all the 4-pizzas at once. This is because the resulting graph still has a well-formed atlas, and therefore is 3-connected, and so the above argument applies for each removed 4-pizza.

5 Advanced Separations

In this section we prove the necessary properties of the separators introduced in Definition 3.10.

5.1 Separating Edges

Definition 5.1 *A* shrinkable segment *S in $\mathcal{M}$ is a (u,v)-segment in $\mathcal{M}$ such that (i) $\{u,v\}$ is an unmarked edge in G, (ii) both endpoints of S are 3-points, and (iii) the endpoints of S are touched by distinct nations a and b which are adjacent in G. Nations a and b are called the* ending nations *of S.*

In the next two results, we show a close relationship between separating edges and shrinkable segments.

Lemma 5.2 *Assume that G has a separating edge $\{a,b\}$. Let $G' = G - \{a,b\} - \mathcal{E}[a,b]$. Then, for every $\{x,y\} \in E$ such that x and y belong to different connected components of G', $\langle a,x,b,y\rangle$ is a correct 4-pizza in G.*

Proof: Let $\mathcal{M}'$ be the atlas of G obtained from $\mathcal{M}$ by contracting those shrinkable segments whose ending nations are a and b. All edges except $\{a,b\}$ are good in $\mathcal{M}'$.

First, we claim that for every $\{u,v\} \in E$ such that u and v belong to different connected components of G', there is a point in $\mathcal{M}'$ at which nations u, a, v, b meet at p cyclically in this order. Towards a contradiction, assume that such a point does not exist in $\mathcal{M}'$. By the definition of G', $\{u,v\}$ is in $\mathcal{E}[a,b]$. There is no nation $w \in V - \{a,b,u,v\}$ adjacent to both u and v; otherwise, w would connect u and v in G' by Fact 3.9. By the fact that $\mathcal{M}'$ has no hole, nations u and v only share a unique curve segment S in $\mathcal{M}'$ and the endpoints of S can be touched only by a or b in $\mathcal{M}'$. Nation a cannot touch both endpoints of S; otherwise, since the edges $\{a,u\}$ and $\{a,v\}$ are still good in $\mathcal{M}'$, nations a, u, and v would have to occupy the whole sphere, a contradiction. Similarly, b cannot touch both endpoints of S. So, both endpoints of S are 3-points. In summary, one endpoints of S is touched by a and the other is touched by b. Then S would be a shrinkable segment whose ending nations are a and b, contradicting the choice of $\mathcal{M}'$.

Second, we claim that there is no (a,b)-segment in $\mathcal{M}'$. Towards a contradiction, assume that an (a,b)-segment S exists in $\mathcal{M}'$. By the first claim, there is an (a,b)-point p in $\mathcal{M}'$. Note that p is not on S. Let x and y be the two nations of $V - \{a,b\}$ that meet at p. Since $\mathcal{M}'$ has no hole, G is 4-connected and $|V| \geq 9$, there is a nation $z \in V - \{a,b,x,y\}$ that touches either x or y in $\mathcal{M}'$. If z touches x (respectively, y) in $\mathcal{M}'$, then z is not reachable from y (respectively, x) in $G - \{a,b,x\}$ (respectively, $G - \{a,b,y\}$), contradicting Assumption 4.

Third, we claim that there are at least two (a,b)-points in $\mathcal{M}'$. Otherwise, if there is only one (a,b)-point, then by the first claim, $\mathcal{E}[a,b]$ would have at most one edge, and the 4-connectivity of G would prevent the separation of G'.

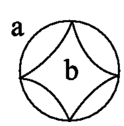

Figure 5.1: Layout $\mathcal{M}'|_{\{a,b\}}$ when $\ell = 4$.

Let ℓ be the number of (a,b)-points. Since $\ell \geq 2$ and there is no (a,b)-segment, we see that the atlas $\mathcal{M}'$ has a cyclic sequence of (a,b)-points $q_0, \ldots, q_{\ell-1}$. These points alternate with ℓ 2-holes in $\mathcal{M}'|_{\{a,b\}}$; Figure 5.1 displays $\mathcal{M}'|_{\{a,b\}}$ when $\ell = 4$.

For each j, let x_j and y_j be the nations meeting a and b. We claim that $\{a, b, x_j, y_j\}$ is an MC_4; otherwise to form a containing 5-clique would force $\ell \leq 3$ and $\mathcal{E}[a, b] = \emptyset$, contradicting the separation of G'. So in fact each q_j corresponds to an edge (x_j, y_j) in $\mathcal{E}[a, b]$, and the components of G' correspond to the set of nations in each hole.

Now consider a particular edge (x_j, y_j). To show that $\langle a, x_j, b, y_j\rangle$ is a correct 4-pizza, we must find a well-formed atlas of G where they meet as they do in $\mathcal{M}'$. This is easy to do: we simply erase all (a, b)-points in $\mathcal{M}'$ except q_j, and the resulting atlas is well-formed. □

Corollary 5.3 *Suppose G has an edge $\{a, b\}$, not inside a 5-clique. Then, $\{a, b\}$ is a separating edge if and only if there is a shrinkable segment in $\mathcal{M}$ with ending nations a and b.*

Proof: The "if" part is obvious from the proof of Lemma 5.2. To prove the "only if" part, suppose there is a shrinkable segment S in $\mathcal{M}$ with ending nations a and b. Let $\mathcal{M}'$ be the atlas of G obtained from $\mathcal{M}$ by contracting S to a single point p. Similarly to the second claim in the proof of Lemma 5.2, we can claim that there is no (a, b)-segment in $\mathcal{M}'$. Thus, besides p, there is exactly one (a, b)-point q in $\mathcal{M}'$, inherited from $\mathcal{M}$. Now, $\mathcal{M}'|_{\{a,b\}}$ has exactly two holes $\mathcal{H}_0$ and $\mathcal{H}_1$. Let Z_0 (respectively, Z_1) be the set of nations of $V - \{a, b\}$ occupying $\mathcal{H}_0$ (respectively, $\mathcal{H}_1$) in atlas $\mathcal{M}'$. Let $x \in Z_0$ and $y \in Z_1$ be the two nations that meet at p in $\mathcal{M}'$. Similarly, let $x' \in Z_0$ and $y' \in Z_1$ be the two nations that meet at q in $\mathcal{M}'$. By $\mathcal{M}'$, edges $\{x, y\}$ and $\{x', y'\}$ are not marked in G and they are all the edges connecting nations of Z_0 to those of Z_1. Moreover, since no 5-clique contains $\{a, b\}$, $\{a, b, x, y\}$ and $\{a, b, x', y'\}$ are MC_4's of G. Therefore, no connected component of $G - \{a, b\} - \mathcal{E}[a, b]$ contains both the nations of Z_0 and those of Z_1. In other words, $\{a, b\}$ is a separating edge of G. □

5.2 Separating 4-Cycles

Lemma 5.4 *Suppose that G has a separating 4-cycle C with vertices a, b, c, d appearing on $G[C]$ in this order. Then, $G - C$ has exactly two connected components G_1 and G_2; and for each $i \in \{1, 2\}$, the marked graph G'_i obtained from $G[V(G_i) \cup C]$ by adding edge $\{a, c\}$ and marking edges $\{a, b\}$, $\{b, c\}$, $\{c, d\}$, $\{d, a\}$, $\{a, c\}$ has a well-formed atlas. Moreover, given a well-formed atlas of G'_1 and one of G'_2, we can easily construct one of G.*

Proof: Since $G[C]$ is a cycle and $\mathcal{M}$ is well-formed, there are exactly two holes $\mathcal{H}_1$ and $\mathcal{H}_2$ in $\mathcal{M}|_C$. For $j \in \{1, 2\}$, let U_j be the set of nations that occupy $\mathcal{H}_j$ in atlas $\mathcal{M}$. Clearly, the nations in U_j are connected together in $G - C$. By this and the assumption that $G - C$ is disconnected, both $G[U_1]$ and $G[U_2]$ are connected components of $G - C$ and $G - C$ has no other connected component. So, $\mathcal{H}_1$ and $\mathcal{H}_2$ must be disjoint. In turn, for each edge $\{u, v\}$ in $G[C]$, there is a (u, v)-segment in $\mathcal{M}$.

For each $j \in \{1, 2\}$, there is a unique hole in $\mathcal{M}|_{U_j \cup C}$ and it is (strongly) touched only by the nations of C. So, modifying $\mathcal{M}|_{U_j \cup C}$ by extending nation a to cover its unique hole yields a well-formed atlas of G'_j in which nations a, b, and c meet at a 3-point and nations a, c, and d meet at a 3-point. So, by Lemma 4.3, both $G'_j - \{a, b, c\}$ and $G'_j - \{a, c, d\}$ are connected.

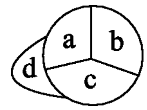

Figure 5.2: Layout $\mathcal{M}_j|_C$.

Suppose that we are given an atlas $\mathcal{M}_j$ for each G'_j. Since $G'_j - \{a, b, c\}$ is connected and the three edges $\{a, b\}$, $\{a, c\}$, and $\{c, b\}$ are marked in G'_j, nations a, b, and c meet at a 3-point in $\mathcal{M}_j$, by Lemma 4.3. Similarly, nations a, c, and d must meet at a 3-point in $\mathcal{M}_j$. Thus, by the well-formedness of $\mathcal{M}_j$, Figure 5.2 displays $\mathcal{M}_j|_C$. By the figure, we can cut off a small area inside $\mathcal{M}_j$ around the (a, c)-segment; the nations in the other atlas can be embedded in the resulting open area. □

5.3 Separating Triples

Lemma 5.5 *Suppose G has no separating edge but has a separating triple $\langle a, b, c\rangle$. Then, $G - \{a, b, c\} - \mathcal{E}[a, b]$ has exactly two connected components G_1 and G_2. Moreover, $\langle a, u, b, v\rangle$ is a correct 4-pizza, where $\{u\} = V(G_1) \cap N_G(V(G_2))$ and $\{v\} = V(G_2) \cap N_G(V(G_1))$.*

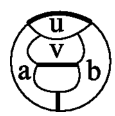

Figure 5.3: A display of $G[\{a, b, u, v\}]$.

Proof: Let $C = \{a, b, c\}$ and $G' = G - C - \mathcal{E}[a, b]$. Since G is 4-connected, $G - C$ is connected. So $\mathcal{E}[a, b]$ is non-empty to disconnect G', and we may choose $\{u, v\} \in \mathcal{E}[a, b]$ such that u and v belong to different components of G'. By definition of $\mathcal{E}[a, b]$, $\{a, b, u, v\}$ is an MC_4 in G.

We claim that nations u and v do not strongly touch in $\mathcal{M}$. Assume, on the contrary, that a (u, v)-segment S exists in $\mathcal{M}$. Since $\mathcal{M}$ has no hole, there are nations w_1, w_2 in $V - \{u, v\}$ such that w_1 touches one endpoint of S and w_2 touches the other. If w_1 were not in $\{a, b\}$, then w_1 would connect u and v in G' in G' by Fact 3.9. Thus $w_1 \in \{a, b\}$, and similarly $w_2 \in \{a, b\}$. By the well-formedness of $\mathcal{M}$ and the fact that $|V| > 3$, we can verify that there is no way for nation a or b to touch both endpoints of S. So, nation a touches one endpoint of S and b touches the other; Figure 5.3 displays $\mathcal{M}|_{\{a,b,u,v\}}$. By Figure 5.3 and the fact that edge $\{u, v\}$ is not marked in G, $\{a, b\}$ is a separating edge of G, a contradiction. Thus, the claim holds.

By the claim, nations u and v weakly touch at a point p in $\mathcal{M}$. Since $\mathcal{M}$ has no hole, there are two distinct nations w_1, w_2 in $V - \{u, v\}$ which meet at p in $\mathcal{M}$. As before, we can show that $\{w_1, w_2\} = \{a, b\}$. Thus, the four nations a, u, b, v appear around p cyclically in this order in $\mathcal{M}$. In turn, by the well-formedness of $\mathcal{M}$, $\langle a, u, b, v\rangle$ is a correct 4-pizza of G.

The discussions above actually prove that for every pair of adjacent nations x and y of G that belong to different connected components of G', nations a, x, b, y must meet at a

4-point cyclically in this order in $\mathcal{M}$. Since p is the unique point in $\mathcal{M}$ at which a and b meet, (u, v) is the unique pair of adjacent nations of G that belong to different connected components of G'. We now claim that the connected components of G' are only G_1 and G_2. Assume, on the contrary, that G' has a connected component G_3 other than G_1 and G_2. Then, there exists a nation $w \in V - (C \cup V(G_3))$ which touches some nation w' of G_3 in $\mathcal{M}$; otherwise, G_3 would be a connected component of $G - C$, a contradiction. But now, (w, w') is a pair of adjacent nations of G that belong to different connected components of G', a contradiction. Thus, the connected components of G' are only G_1 and G_2. We may assume that $u \in V(G_1)$ and $v \in V(G_2)$. G_1 and G_2 only touch at p; otherwise, nations a and b would have to meet at a point other than p in $\mathcal{M}$, a contradiction against the well-formedness of $\mathcal{M}$. Hence, $\{u\} = V(G_1) \cap N_G(V(G_2))$ and $\{v\} = V(G_2) \cap N_G(V(G_1))$. □

5.4 Separating Quadruples

Using Lemma 5.4, we can modify the proof of Lemma 5.5 to prove the following:

Lemma 5.6 *Suppose that G has neither separating edge nor separating 4-cycle, but has a separating quadruple $\langle a, b, c, d\rangle$. Then, $G - \{a, b, c, d\} - \mathcal{E}[a, b]$ has exactly two connected components G_1 and G_2. Moreover, $\langle a, u, b, v\rangle$ is a correct 4-pizza, where $\{u\} = V(G_1) \cap N_G(V(G_2))$ and $\{v\} = V(G_2) \cap N_G(V(G_1))$.*

Corollary 5.7 *Suppose that G does not have an MC_5 or a separating edge. Then, G has a separating quadruple if and only if for some induced 4-cycle C in G, at most one pair of adjacent nations of C weakly touch in $\mathcal{M}$.*

Proof: The "if" part is obvious. For the "only if" part, suppose G has a separating quadruple Q. If G has a separating 4-cycle C, then as observed in the proof of Lemma 5.4, each pair of adjacent nations of C strongly touch in $\mathcal{M}$. Otherwise, as observed in the modified proof of Lemma 5.5 for Lemma 5.6, exactly one pair of adjacent nations of Q weakly touch in $\mathcal{M}$. □

5.5 Separating Triangles

The previous results of this section allow our algorithm to simplify G whenever it contains a separating edge, triple, or quadruple; so, for the rest of this paper we assume:

Assumption 5 *G does not have a separating edge, triple, or quadruple.*

Suppose G has a separating triangle $\langle a, b, c\rangle$. Let C and G' be as described in Definition 3.10(5). Our goal is to show that using C and G', our algorithm can proceed by finding correct 4-pizzas. We begin with three preliminary claims.

Claim 5.8 *If $\{u, v\}$ is an edge in $G - C$ but not in G', then $a \in N_G(u) \cap N_G(v)$. Also, nations u, v, b, and c cannot meet at a 4-point in $\mathcal{M}$.*

Proof: Since $\{u, v\} \in \mathcal{E}[a, b] \cup \mathcal{E}[a, c]$, either $\{a, b, u, v\}$ or $\{a, c, u, v\}$ is an MC_4 of G. In both cases, $a \in N_G(u) \cap N_G(v)$. For the last part, such a 4-point would imply a 5-clique containing the MC_4, contradicting its maximality. □

Claim 5.9 *For every connected component K of G', (i) $C \subseteq N_G(V(K))$ and (ii) G' has another connected component J such that $V(K) \cap N_G(V(J)) \neq \emptyset$.*

Proof: For (i), let $S = C \cap N_G(V(K))$. Since $G - C$ is connected, some edge in $\mathcal{E}[a, b] \cup \mathcal{E}[a, c]$ connects K to an outside vertex, so to support the corresponding MC_4, S must contain either $\{a, b\}$ or $\{a, c\}$. If $|S| = 2$, then S would be a separating edge of G, separating K from the rest. So, $S = C$.

For (ii), if on the contrary $V(K) \cap N_G(V(J)) = \emptyset$ for every J, then K would be a component of $G - C$, contradicting Assumption 4. □

Claim 5.10 *Let Z be a subset of $V - C$. Suppose that a subset $\{u, v, w\}$ of Z is a triangle of G such that u and v belong to different connected components of $G'[Z]$. Then, the following hold:*

1. *Either (i) $C \subseteq N_G(u)$ and $\{C \cap N_G(v), C \cap N_G(w)\} = \{\{a, b\}, \{a, c\}\}$ or (ii) $C \subseteq N_G(v)$ and $\{C \cap N_G(u), C \cap N_G(w)\} = \{\{a, b\}, \{a, c\}\}$.*
2. *There is no $x \in Z - \{u, v, w\}$ with $\{u, v, w\} \subseteq N_G(x)$.*

Note that we typically use this claim with $Z = V - C$.

Proof: Since u and v are disconnected in $G'[Z]$, at least two of the triangle edges are removed. Claim 5.8 applied to these edges implies $\{u, v, w\} \subseteq N_G(a)$. On the other hand, by Fact 3.9 at most one triangle edge is in each of $\mathcal{E}[a, b]$ and $\mathcal{E}[a, c]$, so in fact exactly two edges are removed, and either edge $\{u, w\}$ or $\{v, w\}$ remains in G'.

We suppose $\{v, w\}$ remains, the other case is similar (swap u and v). We also suppose $\{u, v\} \in \mathcal{E}[a, b]$ and $\{u, w\} \in \mathcal{E}[a, c]$, the other case is similar (swap b and c). Then $\{a, b, c\} \subseteq N_G(u)$, $\{a, b\} \subseteq N_G(v)$, and $\{a, c\} \subseteq N_G(w)$. On the other hand, G cannot have the edges $\{v, c\}$ or $\{w, b\}$, since either would imply a 5-clique containing an MC_4. So, the first assertion holds.

For the second assertion, suppose on the contrary there is an $x \in Z - \{u, v, w\}$ with $\{u, v, w\} \subseteq N_G(x)$. As above, we suppose that both $\{a, b, u, v\}$ and $\{a, c, u, w\}$ are MC_4's of G. Then neither $\{a, b\}$ nor $\{a, c\}$ is a subset of $N_G(x)$, since otherwise x would extend one of these MC_4's to a 5-clique. But then the edges from x would all survive in $G'[Z]$, contradicting the disconnection of u and v. □

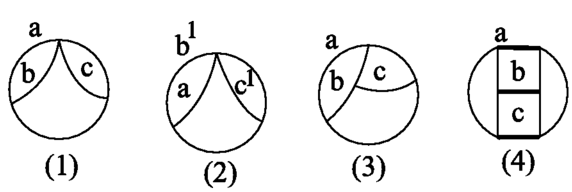

Figure 5.4: Possible displays of a separating triangle $\langle a, b, c\rangle$.

Note that $\mathcal{M}|_C$ can have at most two holes. If $\mathcal{M}|_C$ has only one hole, then Figure 5.4(1), (2), or (3) displays it; otherwise, Figure 5.4(4) displays it. In the next three lemmas, we will show that in fact only Figure 5.4(4) is possible.

Lemma 5.11 *Figure 5.4(1) does not display $\mathcal{M}|_C$.*

Proof: Assume, on the contrary, that Figure 5.4(1) displays $\mathcal{M}|_C$. Let p be the point in $\mathcal{M}|_C$ at which nations a, b, and c meet. Let $p_{a,b}$ (respectively, $p_{a,c}$) be the endpoint of the (a,b)-segment (respectively, (a,c)-segment) other than p in $\mathcal{M}$. There must exist a nation $d \in V - C$ which touches p in $\mathcal{M}$. By the well-formedness of $\mathcal{M}$, nation d touches a only at p and $\{a,d\}$ is not a marked edge in G. Let G'_d be the connected component of G' containing d. Let K be a connected component of G' other than G'_d such that some nation u of G'_d touches some nation v of K in $\mathcal{M}$; K exists by Claim 5.9.

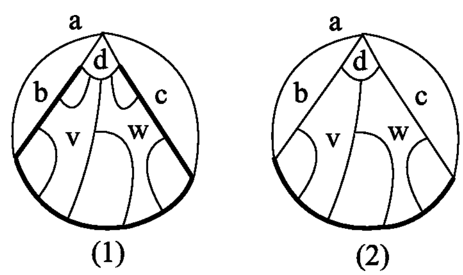

Figure 5.5: Possible displays of $G[\{a,b,c,d,v,w\}]$.

We claim that nation a touches some nation of $G'_d - \{d\}$ in $\mathcal{M}$. Assume, on the contrary, that the claim is false. Clearly, $\{a,b,u,v\}$ or $\{a,c,u,v\}$ is an MC$_4$ of G. Since no nation of $G'_d - \{d\}$ touches a in $\mathcal{M}$, $u = d$. That is, $\{a,b,d,v\}$ or $\{a,c,d,v\}$ is an MC$_4$ of G. Since $\{a,b,c\} \subseteq N_G(d)$, we have $|N_G(v) \cap \{b,c\}| = 1$; otherwise, $\{a,b,c,d,v\}$ would be a 5-clique of G. We assume that $N_G(v) \cap \{b,c\} = \{b\}$; the other case is similar. Then, since nation v cannot touch nation c in $\mathcal{M}$ and $\mathcal{M}$ has no hole, there is a point in $\mathcal{M}$ at which nations v, d and some $w \in V - \{a,b,c,d,v\}$ meet. By Claim 5.10, $C \cap N_G(w) = \{a,c\}$ and there is no $x \in V - \{a,b,c,d,v,w\}$ such that $\{d,v,w\} \subseteq N_G(x)$. Now, we see that Figure 5.5(1) displays $\mathcal{M}|_{\{a,b,c,d,v,w\}}$. There is no $x \in V - \{a,b,c,d,v,w\}$ with $\{d,v\} \subseteq N_G(x)$; otherwise, $C \cap N_G(x) = \{a,c\}$ by Claim 5.10(1), which is impossible by Figure 5.5(1). Similarly, there is no $x \in V - \{a,b,c,d,v,w\}$ with $\{d,w\} \subseteq N_G(x)$. Thus, Figure 5.5(1) is transformable to Figure 5.5(2). By Figure 5.5(2) and the fact that $\{d,a\}$ is not a marked edge in G, $\langle b,c,w,v \rangle$ is a separating quadruple of G, a contradiction. So, the claim holds.

Next, we claim that for every connected component K' of G', there is no point q in $\mathcal{M}$ at which two nations x and y of K' together with two nations w and z of $V - V(K')$ meet cyclically in the order x, w, y, z. Assume, on the contrary, that such q exists in $\mathcal{M}$. Then, by Claim 5.10(2), $C \cap \{w,z\} \neq \emptyset$. By Figure 5.4(1), $q \notin \{p, p_{a,b}, p_{a,c}\}$ and hence $|C \cap \{w,z\}| \leq 1$. So, $|C \cap \{w,z\}| = 1$. In turn, $C \cap \{w,z\} = \{a\}$; otherwise, by Claim 5.8, $\{x,y,a,w,z\}$ would be a 5-clique of G, a contradiction. We assume that $w = a$; the other case is similar. Now, by Claim 5.10(1), $\{C \cap N_G(x), C \cap N_G(y)\} = \{\{a,b\},\{a,c\}\}$ and $C \subseteq N_G(z)$. We assume that $C \cap N_G(x) = \{a,b\}$ and $C \cap N_G(y) = \{a,c\}$; the other case is similar. In summary, Figure 5.6(1) displays $\mathcal{M}|_{\{a,b,c,x,y,z\}}$. There is no $f \in V - \{a,b,c,x,y,z\}$ with $\{x,z\} \subseteq N_G(f)$; otherwise, by Claim 5.10(1), $C \cap N_G(f) = \{a,c\}$

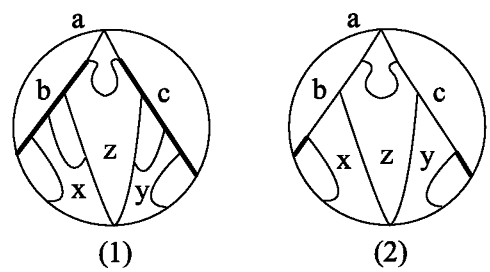

Figure 5.6: Possible displays of $G[\{a, b, c, x, y, z\}]$.

which is impossible by Figure 5.6(1). Similarly, there is no $f \in V - \{a, b, c, x, y, z\}$ with $\{y, z\} \subseteq N_G(f)$. So, Figure 5.6(1) is transformable to Figure 5.6(2). By the latter figure, $\langle a, z, b\rangle$ would be a separating triple of G, a contradiction. So, the claim holds.

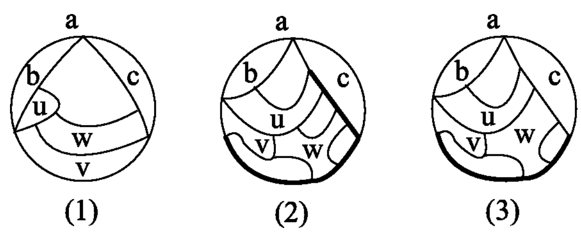

Figure 5.7: Possible displays of $G[\{a, b, c, u, v, w\}]$.

By the above two claims and the fact (Claim 5.9) that $\{a, b, c\} \subseteq N_G(V(K))$, $p_{a,b}$ or $p_{a,c}$ is touched by both G'_d and K in $\mathcal{M}$. Suppose that $p_{a,b}$ is touched by both G'_d and K; the other case is similar. Let u (respectively, v) be the nation of G'_d (respectively, K) touching $p_{a,b}$. Then, the boundaries of nations u and v in $\mathcal{M}$ share a curve segment S. One endpoint of S is $p_{a,b}$. Let q be the other endpoint of S. Neither nation d nor c touches q in $\mathcal{M}$; otherwise, $\{u, v, a, b, d\}$ or $\{u, v, a, b, c\}$ would be a 5-clique of G. By the well-formedness of $\mathcal{M}$, it is impossible that nation a or b touches q. In turn, since $\mathcal{M}$ has no hole, there is a nation $w \in V - \{a, b, c, d, u, v\}$ that touches q in $\mathcal{M}$. Now, by Claim 5.10(1), $C \cap N_G(w) = \{a, c\}$ and either (i) $C \subseteq N_G(v)$ and $C \cap N_G(u) = \{a, b\}$ or (ii) $C \subseteq N_G(u)$ and $C \cap N_G(v) = \{a, b\}$. In case (i) holds, Figure 5.7(1) displays $\mathcal{M}|_{\{a,b,c,u,v,w\}}$ and $\langle b, c, w, u\rangle$ would be a separating quadruple, a contradiction. So, (ii) holds and only Figure 5.7(2) can possibly display $\mathcal{M}|_{\{a,b,c,u,v,w\}}$. There is no $f \in V - \{a, b, c, u, v, w\}$ with $\{u, w\} \subseteq N_G(f)$; otherwise, by Claim 5.10(1), $C \cap N_G(f) = \{a, b\}$ which is impossible by Figure 5.7(2). By this, Figure 5.7(2) is transformable to Figure 5.7(3), by which $\langle a, u, c\rangle$ would be a separating triple of G, a contradiction. This completes the proof. □

Lemma 5.12 *Figure 5.4(2) does not display $\mathcal{M}|_C$.*

Proof: Assume, on the contrary, that Figure 5.4(2) displays $\mathcal{M}|_C$. We assume that $\langle b^1, c^1\rangle = \langle b, c\rangle$ in the figure; the other case is similar. Define points p and $p_{a,b}$, nation d and G'_d as in the proof of Lemma 5.11. By the well-formedness of $\mathcal{M}$, nation d meets b only at p and $\{b, d\}$ is not a marked edge in G. Let $p_{b,c}$ be the endpoint of the (b, c)-segment other than p in $\mathcal{M}$.

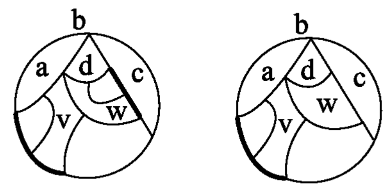

Figure 5.8: Possible displays of $G[\{a, b, c, d, v, w\}]$.

We claim that $G'_d - \{d\}$ touches nation b in $\mathcal{M}$. Assume, on the contrary, that $G'_d - \{d\}$ does not touch nation b. Let K be a connected component of G' other than G'_d such that some nation u of G'_d touches some nation v of K in $\mathcal{M}$. By Claim 5.9, such K exists. Clearly, $\{a, b, u, v\}$ or $\{a, c, u, v\}$ is an MC_4 of G.

Case 1: $u \neq d$. Then $C \cap N_G(u) = \{a, c\}$ and $\{a, c, u, v\}$ is an MC_4 of G. Moreover, there is no $w \in V - \{a, b, c, u, v\}$ with $\{u, v\} \subseteq N_G(w)$; otherwise, since $C \cap N_G(u) = \{a, c\}$, we would have $C \subseteq N_G(v)$ and $C \cap N_G(w) = \{a, b\}$ by Claim 5.10(1), and in turn w would be a vertex of $G'_d - \{d\}$ that touches nation b in $\mathcal{M}$, a contradiction. So, by the fact that $\mathcal{M}$ has no hole, the boundaries of nations u and v share a curve segment S in $\mathcal{M}$, and both endpoints of S are 3-points one of which is touched by a and the other is touched by c in $\mathcal{M}$. By this, S is a shrinkable segment in $\mathcal{M}$, u and v fall into different connected components of $G - \{a, c\} - \mathcal{E}[a, c]$, and $\{a, c\}$ would be a separating edge of G, a contradiction.

Case 2: $u = d$. Then $\{a, b, d, v\}$ or $\{a, c, d, v\}$ is an MC_4 of G. Since $\{a, b, c\} \subseteq N_G(d)$, we have $|N_G(v) \cap \{b, c\}| = 1$; otherwise, $\{a, b, c, d, v\}$ would be a 5-clique of G. So we have two sub-cases.

Case 2.1: $N_G(v) \cap \{b, c\} = \{b\}$. Then $C \cap N_G(v) = \{a, b\}$ and $\{a, b, d, v\}$ is an MC_4 of G. Moreover, since nation v cannot touch nation c in $\mathcal{M}$ and $\mathcal{M}$ has no hole, there is a point in $\mathcal{M}$ at which nations v, d and some $w \in V - \{a, b, c, d, v\}$ meet. By Claim 5.10(1), $C \cap N_G(w) = \{a, c\}$ and there is no $x \in V - \{a, b, c, d, v, w\}$ such that $\{d, v, w\} \subseteq N_G(x)$. Now, we see that Figure 5.8(1) displays $\mathcal{M}|_{\{a,b,c,d,v,w\}}$. There is no $x \in V - \{a, b, c, d, v, w\}$ with $\{d, w\} \subseteq N_G(x)$; otherwise, $C \cap N_G(x) = \{a, b\}$ by Claim 5.10(1), which is impossible by Figure 5.8(1). Thus, Figure 5.8(1) is transformable to Figure 5.8(2). By Figure 5.8(2), $\langle b, v, w, c\rangle$ is a separating quadruple of G, a contradiction.

Case 2.2: $N_G(v) \cap \{b, c\} = \{c\}$. If there is a $w \in V - \{a, b, c, d, v\}$ with $\{d, v\} \subseteq N_G(w)$, then similarly to Case 2.1, we can prove that $\langle b, w, v, c\rangle$ would be a separating quadruple of G, a contradiction. Otherwise, the boundaries of nations d and v share a curve segment S in $\mathcal{M}$, and both endpoints of S are 3-points one of which is touched by a and the other is touched by c in $\mathcal{M}$; by this, S is a shrinkable segment in $\mathcal{M}$, d and v fall into different connected components of $G - \{a, c\} - \mathcal{E}[a, c]$, and $\{a, c\}$ would be a separating edge of G, a contradiction.

Therefore, the claim holds: $G'_d - \{d\}$ touches b.

Next, we claim that for every connected component K' of G', there is no point q in $\mathcal{M}$ at which two nations x and y of K' together with two nations w and z of $V - V(K')$ meet cyclically in the order x, w, y, z. Assume, on the contrary, that such q exists in $\mathcal{M}$. Then, by Claim 5.10(2), $C \cap \{w, z\} \neq \emptyset$. By Figure 5.4(2), $q \notin \{p, p_{a,b}, p_{b,c}\}$ and hence $|C \cap \{w, z\}| \leq 1$. So, $|C \cap \{w, z\}| = 1$. In turn, $C \cap \{w, z\} = \{a\}$; otherwise, by Claim 5.8, $\{x, y, a, w, z\}$

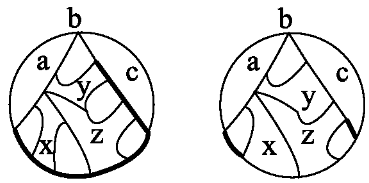

Figure 5.9: Possible displays of $G[\{a, b, c, x, y, z\}]$.

would be a 5-clique of G, a contradiction. We assume that $w = a$; the other case is similar. Now, by Claim 5.10(1), $\{C \cap N_G(x), C \cap N_G(y)\} = \{\{a, b\}, \{a, c\}\}$ and $C \subseteq N_G(z)$. We assume that $C \cap N_G(x) = \{a, b\}$ and $C \cap N_G(y) = \{a, c\}$; the other case is similar. In summary, Figure 5.9(1) displays $\mathcal{M}|_{\{a,b,c,x,y,z\}}$. There is no $f \in V - \{a, b, c, x, y, z\}$ with $\{x, z\} \subseteq N_G(f)$; otherwise, by Claim 5.10(1), $C \cap N_G(f) = \{a, c\}$ which is impossible by Figure 5.9(1). Similarly, there is no $f \in V - \{a, b, c, x, y, z\}$ with $\{y, z\} \subseteq N_G(f)$. So, Figure 5.9(1) is transformable to Figure 5.9(2). By the latter figure, $\langle a, z, b\rangle$ would be a separating triple of G, a contradiction. So, the claim holds.

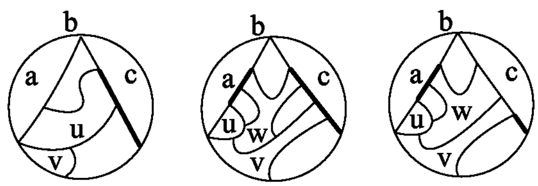

Figure 5.10: Possible displays of $G[\{a, b, c, u, v, w\}]$.

By the above two claims and the fact that $\{a, b, c\} \subseteq N_G(V(K))$, $p_{a,b}$ or $p_{b,c}$ is touched by both G'_d and K in $\mathcal{M}$. By Claim 5.8, $p_{b,c}$ cannot be touched by both G'_d and K. So, $p_{a,b}$ is touched by both G'_d and K. Let u (respectively, v) be the nation of G'_d (respectively, K) touching $p_{a,b}$. Then, the boundaries of nations u and v in $\mathcal{M}$ share a curve segment S. One endpoint of S is $p_{a,b}$. Let q be the other endpoint of S. Neither nation d nor c touches q in $\mathcal{M}$; otherwise, $\{u, v, a, b, d\}$ or $\{u, v, a, b, c\}$ would be a 5-clique of G. By the well-formedness of $\mathcal{M}$, it is impossible that nation a or b touches q. So, there is a nation $w \in V - \{a, b, c, d, u, v\}$ that touches q in $\mathcal{M}$. Now, by Claim 5.10(1), $C \cap N_G(w) = \{a, c\}$ and either (i) $C \subseteq N_G(u)$ and $C \cap N_G(v) = \{a, b\}$ or (ii) $C \subseteq N_G(v)$ and $C \cap N_G(u) = \{a, b\}$. In case (i) holds, Figure 5.10(1) displays $\mathcal{M}|_{\{a,b,c,u,v\}}$; by the figure, it is impossible for nation w to touch all of nations v, a, and c in $\mathcal{M}$, a contradiction. So, only Figure 5.10(2) can possibly display $\mathcal{M}|_{\{a,b,c,u,v,w\}}$. There is no $f \in V - \{a, b, c, u, v, w\}$ with $\{v, w\} \subseteq N_G(f)$; otherwise, by Claim 5.10(1), $C \cap N_G(f) = \{a, b\}$ which is impossible by Figure 5.10(2). By this, Figure 5.10(2) is transformable to Figure 5.10(3), by which $\langle b, u, w, c\rangle$ would be a separating quadruple of G, a contradiction. This completes the proof. □

Lemma 5.13 *Figure 5.4(3) does not display* $\mathcal{M}|_C$.

Proof: Assume, on the contrary, that Figure 5.4(3) displays $\mathcal{M}|_C$. Define points p, $p_{a,b}$ and $p_{a,c}$ as in the proof of Lemma 5.11. Let $p_{b,c}$ be the endpoint of the (b, c)-segment other than p in $\mathcal{M}$. Similarly to the proof of Lemma 5.12, we can prove that for every connected component K' of G', there is no point q in $\mathcal{M}$ at which two nations x and y of K' together with two nations w and z of $V - V(K')$ meet cyclically in the order x, w, y, z.

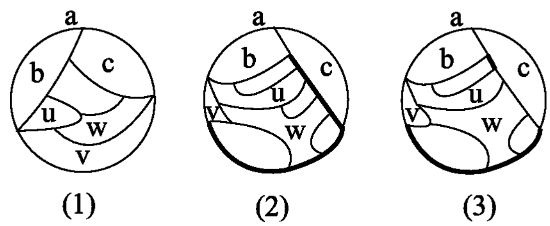

Figure 5.11: Possible displays of $G[\{a, b, c, u, v, w\}]$.

Let G'_1 and G'_2 be two connected components of G' such that $V(G'_1) \cap N_G(V(G'_2)) \neq \emptyset$. By the above claim, Claim 5.8, and Claim 5.9, $p_{a,b}$ or $p_{a,c}$ is touched by both G'_1 and G'_2 in $\mathcal{M}$. We assume that $p_{a,b}$ is touched by both G'_1 and G'_2 in $\mathcal{M}$; the other case is similar. Let u (respectively, v) be the nation of G'_1 (respectively, G'_2) touching $p_{a,b}$. Similarly to the proof of Lemma 5.11, we can prove that there is a nation $w \in V - \{a, b, c, u, v\}$ such that only Figure 5.11(1) or (2) can possibly displays $\mathcal{M}|_{\{a,b,c,u,v,w\}}$. If Figure 5.11(1) displays it, then $\langle b, u, w, c\rangle$ would be a separating quadruple (indeed, a separating 4-cycle) of G, a contradiction. So, suppose that Figure 5.11(2) displays it. Then, there is no $f \in V - \{a, b, c, u, v, w\}$ with $\{u, w\} \subseteq N_G(f)$; otherwise, by Claim 5.10, $C \cap N_G(f) = \{a, b\}$ which is impossible by Figure 5.11(2). By this, Figure 5.11(2) is transformable to Figure 5.11(3). By the latter figure, $\langle b, v, w, c\rangle$ would be a separating quadruple of G, a contradiction. This completes the proof. □

By Lemmas 5.11, 5.12 and 5.13, only Figure 5.4(4) can display $\mathcal{M}|_C$.

Lemma 5.14 *Suppose that $C = \langle a, b, c\rangle$ is a strongly separating triangle of G. Let d be the vertex that constitutes a connected component of G'. Then, there are exactly two vertices $x, y \in V - \{a, b, c, d\}$ such that $\{a, b, d, x\}$ and $\{a, c, d, y\}$ are MC_4 of G. Moreover, both $\langle a, d, b, x\rangle$ and $\langle a, d, c, y\rangle$ are correct 4-pizzas.*

Proof: Let $\mathcal{H}_1$ be one hole of $\mathcal{M}|_C$, and $\mathcal{H}_2$ be the other. Let Z_1 (respectively, Z_2) be the set of nations in $V - C$ that occupy hole $\mathcal{H}_1$ (respectively, $\mathcal{H}_2$) in atlas $\mathcal{M}$. Let $p_{a,b}$ be the point at which nations a and b together with some nation of Z_1 meet in $\mathcal{M}$. Define points $p_{a,c}$ and $p_{b,c}$ similarly.

First, we observe that $C \subseteq N_G(V(K))$ for every connected component K of $G'[Z_1]$. If $K = Z_1$, then this is clear from Figure 5.4(4). Otherwise $G'[Z_1]$ contains some other component K' adjacent to K in $G[Z_1]$, and now our argument resembles that for Claim 5.9(i). That is, let $S = C \cap N_G(V(K))$. Since an edge between K and K' is absent in G', S contains either $\{a, b\}$ or $\{a, c\}$. Assume $S = \{a, b\}$; the $\{a, c\}$ case is similar. Then, in case K is also a connected component of G', it is clear that $\{a, b\}$ would be a separating edge in G, separating K from K'. In case K is not a connected component of G', there is exactly one edge $\{x_1, x_2\} \in E$ with $x_1 \in V(K)$ and $x_2 \in Z_2$; moreover, the four nations a, x_1, b, x_2 must meet at point $p_{a,b}$ in atlas $\mathcal{M}$ cyclically in this order (so, the (a, b)-segment in the layout in Figure 5.4(4) should be contracted to a point). If $\{a, x_1, b, x_2\}$ is an MC_4

of G, then K would be a connected component of G', a contradiction. Otherwise, there is a 5-clique C' containing a, x_1, b, x_2. The nation $x_3 \in C' - \{a, x_1, b, x_2\}$ must belong to Z_1 and touch nation c, in order to touch x_1 and x_2. By this, edge $\{x_1, x_2\}$ remains in G', and $x_3 \in Z_1$, contradicting the fact that K is a connected component of $G'[Z_1]$. So, $S = C$.

Similarly, we have $C \subseteq N_G(V(K))$ for every connected component K of $G'[Z_2]$.

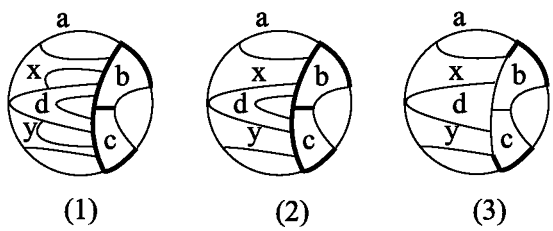

Figure 5.12: Possible displays of $G[\{a, b, c, d, x, y\}]$.

We assume that $d \in Z_1$; the other case is similar. We want to prove that $Z_1 = \{d\}$. Towards a contradiction, assume that $Z_1 \neq \{d\}$. Then, since $\mathcal{M}$ has no hole, there is a connected component K of $G'[Z_1]$ with $V(K) \cap N_G(d) \neq \emptyset$. First, we claim that d and a nation of K must meet at $p_{a,b}$, $p_{a,c}$, or $p_{b,c}$. Assume, on the contrary, that the claim does not hold. Then, since $C \subseteq N_G(V(K)) \cap N_G(d)$ by the above observation, there must exist a point q in $\mathcal{M}$ at which two nations x and y of K together with d and some $u \in C$ meet cyclically in the order x, d, y, u. Claim 5.10 ensures that either (i) $C \cap N_G(x) = \{a, b\}$ and $C \cap N_G(y) = \{a, c\}$ or (ii) $C \cap N_G(x) = \{a, c\}$ and $C \cap N_G(y) = \{a, b\}$. In either case, we have $u = a$. We assume that (i) holds; the other case is similar. Then, Figure 5.12(1) displays $\mathcal{M}|_{\{a,b,c,d,x,y\}}$. There is no $u \in Z_1 - \{d, x, y\}$ with $\{x, d\} \subseteq N_G(u)$; otherwise, by Claim 5.10, $C \cap N_G(u) = \{a, c\}$ which is impossible by Figure 5.12(1). Similarly, there is no $u \in Z_1 - \{d, x, y\}$ with $\{y, d\} \subseteq N_G(u)$. So, Figure 5.12(1) is transformable to Figure 5.12(2). By the latter figure and Claim 5.8, d and each of b and c strongly touch in $\mathcal{M}$. In turn, by the fact that $\langle b, c, d \rangle$ is not a separating triple of G, nations b, c and d meet at a point $q_{b,c,d}$ in $\mathcal{M}$. By Claim 5.8, $q_{b,c,d}$ must be a 3-point; so, b and c strongly touch in $\mathcal{M}$. In summary, Figure 5.12(2) is transformable to Figure 5.12(3). By the latter, $\langle x, y, c, b \rangle$ would be a separating quadruple of G, a contradiction. So, the claim holds: d meets K at $p_{a,b}$, $p_{a,c}$, or $p_{b,c}$.

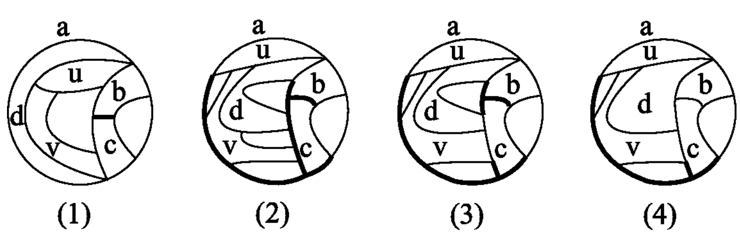

Figure 5.13: Possible displays of $G[\{a, b, c, d, u, v\}]$.

Next, we use the above claim to get a contradiction. By the above claim, d and a nation u of K must meet at $p_{a,b}$, $p_{a,c}$, or $p_{b,c}$ in $\mathcal{M}$. By Claim 5.8, d and u cannot meet at $p_{b,c}$. So, they meet at $p_{a,b}$ or $p_{a,c}$. We assume that they meet at $p_{a,b}$; the other case is similar. Then, the boundaries of nations d and u in $\mathcal{M}$ share a curve segment S. One endpoint of S is $p_{a,b}$. Let q be the other endpoint of S. Nation c cannot touch q in $\mathcal{M}$; otherwise, $\{a, b, c, d, u\}$ would be a 5-clique of G. By the well-formedness of $\mathcal{M}$, it is impossible that only nations a and b touch q. So, there is a nation $v \in V - \{a, b, c, d, u\}$ that touches q in $\mathcal{M}$. Now, by Claim 5.10, $C \cap N_G(v) = \{a, c\}$. Thus, Figure 5.13(1) or (2) displays $\mathcal{M}|_{\{a,b,c,d,u,v\}}$. Actually, the former does not display $\mathcal{M}|_{\{a,b,c,d,u,v\}}$ or else $\langle b, c, a\rangle$ would be a separating triple of G, a contradiction. So, only Figure 5.13(2) can possibly display $\mathcal{M}|_{\{a,b,c,d,u,v\}}$. There is no $w \in Z_1 - \{d, u, v\}$ with $\{d, v\} \subseteq N_G(w)$; otherwise, by Claim 5.10, $C \cap N_G(w) = \{a, b\}$ which is impossible by Figure 5.13(2). By this, Figure 5.13(2) is transformable to Figure 5.13(3). The latter is further transformable to Figure 5.13(4), because each pair of nations in $\{b, c, d\}$ must strongly touch in $\mathcal{M}$ by Claim 5.8 and the fact that $\langle b, c, d\rangle$ is not a separating triple of G. By Figure 5.13(4), $\langle u, b, c, v\rangle$ would be a separating quadruple of G, a contradiction. This completes the proof that $Z_1 = \{d\}$.

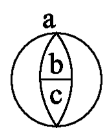

Figure 5.14: An extensible layout of $G[\{a, b, c\}]$.

Now, $Z_1 = \{d\}$. Thus, by Claim 5.8 and Assumption 5 (G has no separating triple), Figure 5.14 displays $\mathcal{M}|_C$. By the figure and the fact that d constitutes a connected component of G', there are exactly two distinct nations x and y of Z_2 such that $\{x, d\} \in E$ and $\{y, d\} \in E$. By the figure, both $\langle a, d, b, x\rangle$ and $\langle a, d, c, y\rangle$ are correct 4-pizzas. Since $Z_2 = V - \{a, b, c, d\}$, finding x and y is easy. This completes the proof of Lemma 5.14. □

Lemma 5.15 *Suppose that there is no strongly separating triangle of G. Further assume that $C = \langle a, b, c\rangle$ is a separating triangle of G. Then, G' has exactly two connected components G_1 and G_2, and there are exactly two edges $\{u, v\}, \{x, y\} \in E$ with $\{u, x\} \subseteq V(G_1)$ and $\{v, y\} \subseteq V(G_2)$. Moreover, $\{a, b, u, v\}$ and $\{a, c, x, y\}$ are MC_4's of G, and both $\langle a, u, b, v\rangle$ and $\langle a, x, c, y\rangle$ are correct 4-pizzas.*

Proof: Define sets Z_1 and Z_2 and points $p_{a,b}$, $p_{a,c}$, and $p_{b,c}$ as in Lemma 5.14. As in the proof of Lemma 5.14, we observe that $C \subseteq N_G(V(K))$ for every connected component K of $G'[Z_1]$ or $G'[Z_2]$.

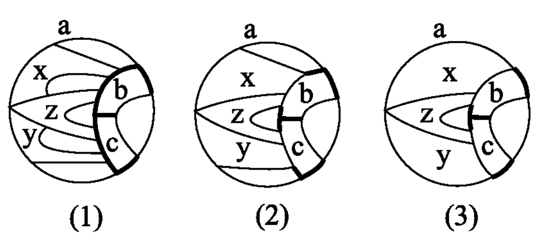

Figure 5.15: Possible displays of $G[\{a, b, c, x, y, z\}]$.

We claim that for every connected component K of $G'[Z_1]$, there is no point q in $\mathcal{M}$ at which two nations x and y of K together with two nations w and z of $(C \cup Z_1) - V(K)$

meet cyclically in the order x, w, y, z. Assume, on the contrary, that such q exists in $\mathcal{M}$. Then, by Claim 5.10 with $Z = Z_1$, $C \cap \{w, z\} \neq \emptyset$. By Figure 5.4(4), $q \notin \{p_{a,b}, p_{a,c}, p_{b,c}\}$ and hence $|C \cap \{w, z\}| \leq 1$. So, $|C \cap \{w, z\}| = 1$. In turn, $C \cap \{w, z\} = \{a\}$; otherwise, by Claim 5.8, $\{x, y, a, w, z\}$ would be a 5-clique of G, a contradiction. We assume that $w = a$; the other case is similar. Now, by Claim 5.10, $\{C \cap N_G(x), C \cap N_G(y)\} = \{\{a, b\}, \{a, c\}\}$ and $C \subseteq N_G(z)$. We assume that $C \cap N_G(x) = \{a, b\}$ and $C \cap N_G(y) = \{a, c\}$; the other case is similar. In summary, Figure 5.15(1) displays $G[\{a, b, c, x, y, z\}]$. There is no $f \in Z_1 - \{x, y, z\}$ with $\{x, z\} \subseteq N_G(f)$; otherwise, by Claim 5.10, $C \cap N_G(f) = \{a, c\}$ which is impossible by Figure 5.15(1). Similarly, there is no $f \in Z_1 - \{x, y, z\}$ with $\{y, z\} \subseteq N_G(f)$. So, Figure 5.15(1) is transformable to Figure 5.15(2). The latter figure is further transformable to Figure 5.15(3), because (i) $\langle a, b, x\rangle$ and $\langle a, c, y\rangle$ are not separating triples of G and (ii) both $\{a, b, x, z\}$ and $\{a, c, y, z\}$ are MC_4's of G. By Figure 5.15(3), $\langle a, b, z\rangle$ would be a strongly separating triangle of G, a contradiction. So, the claim holds.

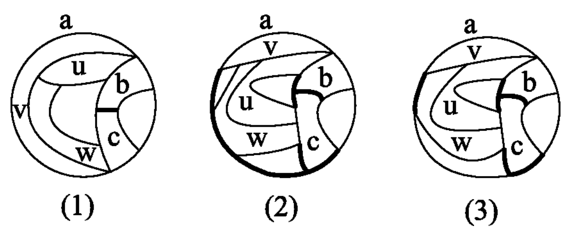

Figure 5.16: Possible displays of $G[\{a, b, c, u, v, w\}]$.

Next, we claim that $G'[Z_1]$ is connected. Assume, on the contrary, that $G'[Z_1]$ is disconnected. Then, since $\mathcal{M}$ has no hole, there are two distinct connected components K and K' of $G'[Z_1]$ such that $V(K) \cap N_G(V(K')) \neq \emptyset$. Since $C \subseteq N_G(V(K))$ and $C \subseteq N_G(V(K'))$, some nation u of K and some nation v of K' have to meet at $p_{a,b}$, $p_{a,c}$ or $p_{b,c}$ in $\mathcal{M}$, by the claim of the previous paragraph and Figure 5.4(4). By Claim 5.8, u and v cannot meet at $p_{b,c}$ in $\mathcal{M}$. We assume that u and v meet at $p_{a,b}$ in $\mathcal{M}$; the other case is similar. Similarly to the proof of Lemma 5.11, we can prove that there is a nation $w \in Z_1 - \{u, v\}$ such that only Figure 5.16(1) or (2) can possibly display $\mathcal{M}|_{\{a,b,c,u,v,w\}}$. Actually, Figure 5.16(1) does not display it or else $\langle b, c, a\rangle$ would be a separating triple of G. So, only Figure 5.16(2) can possibly display it. Since $\langle a, w, v\rangle$ is not a separating triple of G, Figure 5.16(2) is transformable to Figure 5.16(3). By the latter figure, $\langle a, w, u\rangle$ would be a strongly separating triangle of G, a contradiction. So, the claim holds. Similarly, we can prove that $G'[Z_2]$ is connected.

Since both $G'[Z_1]$ and $G'[Z_2]$ are connected, both are connected components of G' and G' has no other connected component. So, by Claim 5.8, the figure obtained from Figure 5.4(4) by contracting the bold (b, c)-segment to a single point does not display $\mathcal{M}|_C$. In turn, the bold (a, b)-segment in Figure 5.4(4) should be contracted to a single point; otherwise, $\langle a, c, b\rangle$ would be a separating triple of G. Similarly,the bold (a, c)-segment in Figure 5.4(4) should be contracted to a single point. Thus, Figure 5.14 displays $\mathcal{M}|_C$. Let $q_{a,b}$ (respectively, $q_{a,c}$) be the point where nations a and b (respectively, nations a and c) meet in $\mathcal{M}|_C$. By the figure, a unique nation $u \in Z_1$ and a unique nation $v \in Z_2$ meet at $q_{a,b}$, and $\{a, b, u, v\}$ is an MC_4 of G. Similarly, a unique nation $x \in Z_1$ and a unique nation $y \in Z_2$ meet at $q_{a,c}$, and $\{a, c, x, y\}$ is an MC_4 of G. Moreover, both $\langle a, u, b, v\rangle$

and $\langle a, x, c, y \rangle$ are correct 4-pizzas. By the figure, other than $\{u, v\}$ and $\{x, y\}$, there is no $\{w_1, w_2\} \in E$ with $w_1 \in Z_1$ and $w_2 \in Z_2$. □

Just for the next corollary, we temporarily drop Assumptions 4 and 5.

Corollary 5.16 *Suppose that G does not have an MC_5, a separating edge, or a separating quadruple. Then, G has a separating triangle if and only if for some 3-clique C of G, (i) the nations of C do not meet at a point in $\mathcal{M}$ and (ii) at least one pair of nations of C strongly touch in $\mathcal{M}$.*

Proof: The "if" part is obvious. For the "only if" part, suppose G has a separating triangle T. If G is not 4-connected, then by Lemma 4.3, there is a 3-clique C in G such that the nations of C do not meet at a point in $\mathcal{M}$ and every pair of nations of C strongly touch in $\mathcal{M}$. So, we may assume that G is 4-connected. Then, by the proof of Lemma 5.5, in case G has a separating triple C', the nations of C' do not meet at a point in $\mathcal{M}$ and at most one pair of nations of C weakly touch in $\mathcal{M}$. Thus, we may further assume that G has no separating triple. Then, by the layouts found in Lemmas 5.11 through 5.15, the nations of T do not meet at a point in $\mathcal{M}$ and at least one pair of nations of T strongly touch in $\mathcal{M}$. □

By the reductions in this section, our algorithm may make progress whenever G has a separating edge, quadruple, or triangle. Hereafter we assume that all such reductions have been made:

Assumption 6 *G does not have a separating edge, quadruple, or triangle.*

In fact Assumption 6 implies Assumptions 4 (by Lemma 4.3(1)) and 5 (by definition). So this one statement summarizes the effect of applying all the reductions in this section and the previous.

6 Removing Maximal 5-Cliques

We assume that G has an MC_5; our goal of this section is to show how to remove MC_5's from G. The idea behind the removal of an MC_5 C from G is to try to find and remove a correct center P of C. By Fact 3.7, we make progress after removing P. After removing P, the resulting G may no longer satisfy Assumption 6; in that case, the algorithm must therefore reapply the reductions of the previous sections before considering another MC_5. Also, not unexpectedly, our search for a correct center of C may fail. In this case, we will be able to decompose G into smaller graphs to make progress.

For a positive integer k, two maximal cliques C_1 and C_2 are *k-sharing* if $|C_1 \cap C_2| = k$.

Every MC_5 C of G is 4-sharing with at most two other MC_5's C' of G; this is because the center of C' must be a 3-point bordering a hole in $\mathcal{M}|_C$, and there are at most two such points in the possible displays of Figure 3.3. We claim that at least one MC_5 of G is 4-sharing with two other MC_5's of G. Towards a contradiction, assume that the claim does not hold. Let $C = \{a, b, c, d, e\}$ be an MC_5 of G. When C is 4-sharing with no MC_5 of G, none of Figure 3.3(1) through (4) displays $\mathcal{M}|_C$ or else either V would equal C or at least one of $\langle e^1, a^1, b^1 \rangle$, $\langle e^1, c^1, d^1 \rangle$, and $\langle e^1, a^1, d^1 \rangle$ would be a separating triangle of G, a contradiction. So, consider the case where C is 4-sharing with exactly one MC_5,

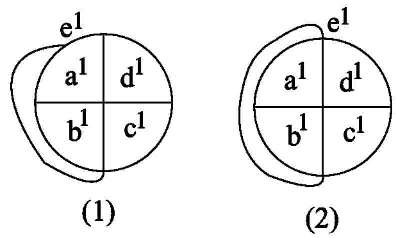

Figure 6.1: Displays of an MC_5 C, 4-sharing with one other.

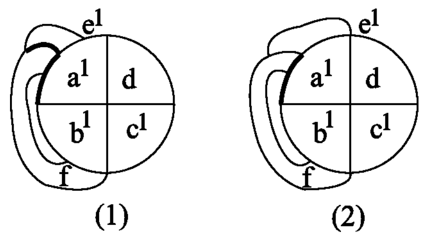

Figure 6.2: Displays of 4-sharing MC_5's C and C_1.

say $C_1 = \{a^1, b^1, c^1, e^1, f\}$, of G. In this case, by Assumption 6 (G has no separating triangle), Figure 3.3(1), (2), and (4) are transformable to Figure 6.1(1), and Figure 3.3(3) is transformable to Figure 6.1(2). By Figure 6.1(1) and (2), only Figure 6.2(1) or (2) can possibly display $\mathcal{M}|_{\{a,\ldots,f\}}$. Actually, Figure 6.2(2) does not display $\mathcal{M}|_{\{a,\ldots,f\}}$; otherwise, since C_1 is 4-sharing with no MC_5 of G other than C, there is no $g \in V - \{a, \ldots, f\}$ with $\{a^1, b^1, e^1, f\} \subseteq N_G(g)$ and $\langle a^1, f, e^1 \rangle$ would be a separating triangle of G, a contradiction. Similarly, Figure 6.2(1) does not display $\mathcal{M}|_{\{a,\ldots,f\}}$; otherwise, since $|V| \geq 9$, $\langle a^1, f, b^1 \rangle$ or $\langle a^1, f, e^1 \rangle$ would be a separating triple of G, a contradiction. Therefore, the claim holds.

By the above claim, if G has an MC_5, then it has an MC_5 that is 4-sharing with two other MC_5's of G. By our assumption that G has an MC_5, G has an MC_5 $C = \{a, b, c, d, e\}$ that is 4-sharing with two other MC_5's, say $C_1 = \{a, c, d, e, f\}$ and $C_2 = \{a, b, c, e, g\}$, of G. Let $U = C \cup \{f, g\}$. We show how to find a correct center of C below. First, we make a simple but useful observation.

Fact 6.1 *Let W be a subset of an MC_5 C' of G with $|W| \geq 3$. If all the edges in $E(G[W])$ are marked in G or $G - C'$ has a vertex x with $W = C' \cap N_G(x)$, then W contains all correct crusts of C'. In particular, if C and C' are MC_5's with $|C \cap C'| \geq 3$, then both crusts are in the intersection.*

Vertices f and g are not adjacent in G; otherwise, only Figure 3.3(3) or (4) can display $\mathcal{M}|_C$, but after drawing nations f and g in the two figures, we see that the 4-connectedness of G would force V to equal U, contradicting Assumption 3. So, only Figure 6.3(1) or Figure 6.3(2) can display $\mathcal{M}|_U$. By the figures, a correct center of C can be found from a correct crust immediately. So, it suffices to find out which one of a, c, and e is a correct crust of C.

Let α be the number of vertices $v \in \{a, c, e\}$ such that $N_G(v) \subseteq U$. $\alpha \leq 1$; otherwise, no matter which of Figure 6.3(1) and (2) displays $\mathcal{M}|_U$, the 4-connectedness of G would force V to equal U, contradicting Assumption 3. First, consider the case where $\alpha = 0$. In

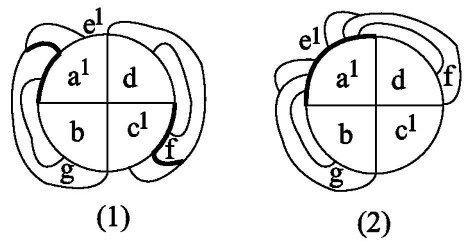

Figure 6.3: Displays of MC_5 C, 4-sharing with C_1 and C_2.

this case, only Figure 6.3(1) displays $\mathcal{M}|_U$. Moreover, by this figure and Assumption 6 (G has no separating triple), there are a unique nation $h \in V - U$ with $\{a^1, b, e^1, g\} \subseteq N_G(h)$. Similarly, there is a unique nation $i \in V - U$ with $\{c^1, d, e^1, f\} \subseteq N_G(i)$. So by Fact 6.1, the unique nation in $N_G(h) \cap N_G(i)$ is a correct crust of C.

Now, we may assume that $\alpha = 1$. We may further assume that c is the unique $u \in \{a, c, e\}$ such that $N_G(u) \subseteq U$. Then, we can delete c from the permutable list $\langle a, c, e \rangle$ in Figure 6.3(2), or more intuitively, we can let $c^1 = c$ in the figure. Similarly, if Figure 6.3(1) displays $\mathcal{M}|_U$, we can let either $a^1 = c$ or $c^1 = c$ in the figure; this would imply either $N_G(\{b, c, g\}) \subseteq U$ or $N_G(\{c, d, f\}) \subseteq U$, respectively. No matter which of Figure 6.3(1) and (2) displays $\mathcal{M}|_U$, if there is a $u \in \{a, e\}$ such that $\{u, d\}$ or $\{u, b\}$ is a marked edge in G, then the unique nation in $\{a, e\} - \{u\}$ is a correct crust of C. So, we may assume that none of $\{a, d\}$, $\{e, d\}$, $\{a, b\}$, and $\{e, b\}$ is a marked edge in G. It remains to consider three cases as follows.

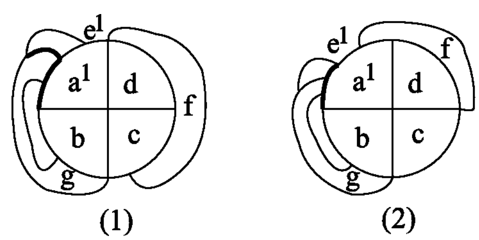

Figure 6.4: Displays of $G[\{a, \ldots, g\}]$ in Case 1.

Case 1: $N_G(\{c, d, f\}) \subseteq U$. Then, Figure 6.3(1) and (2) are transformable to Figure 6.4(1) and (2), respectively.

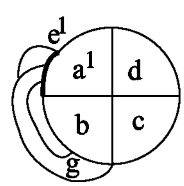

Figure 6.5: A display of $G'[\{a, b, c, d, e, g\}]$ in Case 1.1.

Case 1.1: Edge $\{c, f\}$ is not marked in G. Then, Figure 6.4(1) is transformable to Figure 6.4(2), and hence the latter displays $\mathcal{M}|_U$. Let G' be the marked graph obtained from $G - \{f\}$ by marking the following edges: $\{b, c\}$, $\{c, d\}$, $\{a, e\}$, $\{a, d\}$, $\{e, d\}$. By

Figure 6.4(2), we can obtain a well-formed atlas $\mathcal{M}'$ of G' from $\mathcal{M}$ by extending nation e^1 to occupy nation f. Figure 6.5 displays $\mathcal{M}'|_{\{a,\ldots,e,g\}}$. On the other hand, we claim that every well-formed atlas $\mathcal{M}''$ of G' can be used to construct a well-formed atlas of G. To see this, first note that by Fact 6.1, the crust of C in $\mathcal{M}''$ must be either a or e. Suppose that the crust is e; the other case is similar. Then, since edges $\{b,c\}$ and $\{c,d\}$ are marked in G', the center of C in $\mathcal{M}''$ must be $\langle a,b,c,d\rangle$. Moreover, since $N_{G'}(\{d\}) \subseteq C$, the four nations a, c, d, and e must be related in $\mathcal{M}''$ as shown in Figure 6.5. Thus, we can assign a suitable sub-region of e to f to obtain an atlas of G. This establishes the claim.

Case 1.2: Edge $\{c,f\}$ is marked in G. Then, only Figure 6.4(1) displays $\mathcal{M}|_U$. By the figure, at most one of edges $\{a,f\}$ and $\{e,f\}$ is marked in G. Moreover, if $\{a,f\}$ is marked in G, then a is a correct crust of C. Similarly, if $\{e,f\}$ is marked in G, then e is a correct crust of C. So, it remains to consider the case where neither $\{a,f\}$ nor $\{e,f\}$ is a marked edge in G. In this case, it suffices to construct a marked graph G' as in Case 1.1.

Case 2: $N_G(\{b,c,g\}) \subseteq U$. Similar to Case 1, after relabeling.

Case 3: Neither $N_G(\{b,c,g\}) \subseteq U$ nor $N_G(\{c,d,f\}) \subseteq U$. Then as argued above, Figure 6.3(2) displays $G[U]$. We consider three sub-cases as follows:

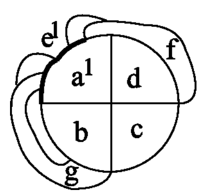

Figure 6.6: A display of $G[\{a,\ldots,g\}]$ in Case 3.1.

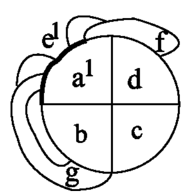

Figure 6.7: A display of $G'[\{a,\ldots,g\}]$ in Case 3.1.

Case 3.1: There is no $v \in V-U$ such that $d \in N_G(v)$ and $N_G(v) \cap \{a,e\} \neq \emptyset$. Then, Figure 6.6 displays $\mathcal{M}|_U$. Let G' be the marked graph obtained from $G - \{\{c,f\}\}$ by marking the following edges: $\{b,c\}$, $\{c,d\}$, $\{a,d\}$, $\{e,d\}$, $\{a,f\}$, $\{e,f\}$, $\{d,f\}$. By Figure 6.6, we can obtain a well-formed atlas $\mathcal{M}'$ of G' by erasing the (c,f)-point in $\mathcal{M}$. Figure 6.7 displays $\mathcal{M}'|_{\{a,\ldots,g\}}$. By Figure 6.7 and Lemma 4.3, both $G' - \{a,d,f\}$ and $G' - \{e,d,f\}$ are connected. We claim that every well-formed atlas $\mathcal{M}''$ of G' can be used to construct a well-formed atlas of G. To see this, first note that by Fact 6.1, the crust of C in $\mathcal{M}''$ must be either a or e. We assume that the crust is e; the other case is similar. Then, since $\{b,c\}$ and $\{c,d\}$ are marked edges in G', the center of C in $\mathcal{M}''$ must be $\langle a,b,c,b\rangle$. Moreover, since $G' - \{a,d,f\}$ is connected, the marked edges $\{a,d\}$, $\{d,f\}$ and $\{f,a\}$ of G' force nations a, d and f to meet at a 3-point in $\mathcal{M}''$. For a similar reason, nations e, d and f meet at a 3-point in $\mathcal{M}''$. Now, since $N_{G'}(c) \subseteq C$, the four nations c, d, e, and f

must be related in $\mathcal{M}''$ as shown in Figure 6.7. Thus, we can assign a suitable sub-region of e to f to obtain a well-formed atlas of G.

Case 3.2: There is no $v \in V - U$ such that $b \in N_G(v)$ and $N_G(v) \cap \{a, e\} \neq \emptyset$. Similar to Case 3.1.

Case 3.3: There are nations h and i in $V - U$ such that $d \in N_G(h)$, $N_G(h) \cap \{a, e\} \neq \emptyset$, $b \in N_G(i)$, and $N_G(i) \cap \{a, e\} \neq \emptyset$. By Figure 6.3(2), no nation of $V - U$ can touch both b and d in $\mathcal{M}$. So, h and i are distinct nations. Moreover, if $|N_G(h) \cap \{a, e\}| = 1$ (respectively, $|N_G(i) \cap \{a, e\}| = 1$), then the unique nation in $\{a, e\} - N_G(h)$ (respectively, $\{a, e\} - N_G(i)$) must be a correct crust and we are done. So, we assume that $\{a, e\} \subseteq N_G(h)$ and $\{a, e\} \subseteq N_G(i)$. Then, by Figure 6.3(2), $\{a, d, e, f, h\}$ and $\{a, b, e, g, i\}$ are MC_5's in G. Let $U_h = U \cup \{h\}$. If $\{g, h\}$ were an edge in G, then by Figure 6.3(2), after drawing nation h in $\mathcal{M}|_U$, we see that the 4-connectedness of G would force V to equal U_h, contradicting Assumption 3. So, $\{g, h\} \notin E$. Similarly, $\{f, i\} \notin E$. Then, Figure 6.8(1) or Figure 6.8(2) displays $\mathcal{M}|_{U_h}$. If edge $\{d, h\}$ is marked in G or $N_G(d) - U_h \neq \emptyset$, Figure 6.8(2) displays $\mathcal{M}|_{U_h}$; otherwise, Figure 6.8(2) is transformable to Figure 6.8(1). So, we can decide whether Figure 6.8(1) or (2) displays $\mathcal{M}|_{U_h}$.

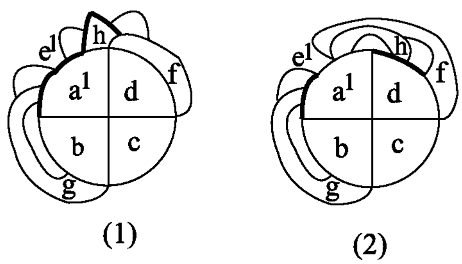

Figure 6.8: Displays of $\mathcal{M}|_{U_h}$ in Cases 3.3.1 and 3.3.2.

Case 3.3.1: Figure 6.8(1) displays $\mathcal{M}|_{U_h}$. We further distinguish two cases as follows.

Case 3.3.1.1: There is no $v \in V - U_h$ such that $f \in N_G(v)$ and $\{a, e\} \cap N_G(v) \neq \emptyset$. Let $\mathcal{D}$ be the figure obtained from Figure 6.8(1) by extending nation h to occupy the hole touched by e^1, f and h. By Figure 6.8(1), $\mathcal{D}$ displays $\mathcal{M}|_{U_h}$ and so $N_G(f) \subseteq U_h$. Moreover, by figure $\mathcal{D}$, if there is a $w \in \{a, e\}$ such that edge $\{w, f\}$ is marked in G, then w is a correct crust of C. So, we may assume that none of the edges $\{a, f\}$ and $\{e, f\}$ is marked in G. Let G' be the marked graph obtained from $G - \{f\}$ by marking the following edges: $\{b, c\}$, $\{c, d\}$, $\{a, d\}$, $\{e, d\}$, $\{a, h\}$, $\{e, h\}$, $\{d, h\}$. By figure $\mathcal{D}$, we can obtain a well-formed atlas $\mathcal{M}'$ of G' from $\mathcal{M}$ by (i) erasing the (c, f)-point and further (ii) extending nation h to occupy f. Indeed, by renaming nation f in Figure 6.7 as h, we obtain a figure displaying $\mathcal{M}'|_{\{a,\ldots,e,g,h\}}$. Moreover, similarly to Case 3.1, we can prove that every well-formed atlas of G' can be used to construct a well-formed atlas of G.

Case 3.3.1.2: There is a $j \in V - U_h$ such that $f \in N_G(j)$ and $\{a, e\} \cap N_G(j) \neq \emptyset$. If $\{a, e\} \not\subseteq N_G(j)$, then by Figure 6.8(1), the unique nation in $\{a, e\} \cap N_G(j)$ is a correct crust of C and we are done. So, we assume that $\{a, e\} \subseteq N_G(j)$. Recall that $\{f, i\} \notin E$. So, $j \neq i$. By Figure 6.8(1), if there is a $w \in \{a, e\}$ such that $\{w, c\}$ is a marked edge in G, then w is a correct crust of C. So, we may assume that neither $\{a, c\}$ nor $\{e, c\}$ is a marked edge in G. Let G' be the graph obtained from $G - \{c, d\}$ by adding the three edges $\{g, f\}$, $\{b, f\}$,

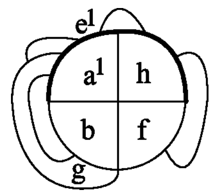

Figure 6.9: Display of $G'[\{a, b, e, f, g, h\}]$ in Case 3.3.1.2.

and $\{h, b\}$ and further marking the two edges $\{b, f\}$ and $\{f, h\}$. By Figure 6.8(1), we can obtain a well-formed atlas $\mathcal{M}'$ of G' from $\mathcal{M}$ by (i) erasing the (d, e^1)-point, (ii) erasing the (a^1, f)-point, (iii) extending nation f to occupy nation c, and (iv) extending nation h to occupy nation d. Indeed, Figure 6.9 displays $\mathcal{M}'|_{\{a,e,b,f,g,h\}}$. We claim that every well-formed atlas $\mathcal{M}''$ of G' can be used to construct a well-formed atlas of G. To see this, first note that G' contains the MC_5's $C' = \{a, e, b, f, h\}$, $C'_1 = \{a, e, b, f, g\}$, $C'_2 = \{a, e, f, h, j\}$, and $C'_3 = \{a, e, b, g, i\}$. These MC_5's and Fact 6.1 ensure that the crust of C' in $\mathcal{M}''$ must be a or e and that the two nations b and h do not appear consecutively around the center of C' in $\mathcal{M}''$. We assume that the crust of C' in $\mathcal{M}''$ is e; the other case is similar. Then, the center of C' in $\mathcal{M}''$ is $\langle a, b, f, h\rangle$. This together with the well-formedness of $\mathcal{M}''$ implies that the crust of C'_1 in $\mathcal{M}''$ is either a or f. On the other hand, since $C'_1 \cap N_G(i) = \{a, e, b, g\}$, f is not a correct crust of C'_1 by Fact 6.1. Thus, the crust of C'_1 in $\mathcal{M}''$ is a. Therefore, the centers of C' and C'_1 are as shown in Figure 6.9. From this, the claim follows immediately.

Case 3.3.2: Figure 6.8(2) displays $\mathcal{M}|_{U_h}$. In this case, we check if there is a $v \in V - U_h$ such that $d \in N_G(v)$ and $N_G(v) \cap \{a, e\} \neq \emptyset$. If such v exists, then $|N_G(v) \cap \{a, e\}| = 1$ and the unique nation in $\{a, e\} - N_G(v)$ is a correct crust of C. If no such v exists, then by Figure 6.8(2) and the 4-connectedness of G, we have $N_G(\{d, f, h\}) \subseteq U_h$ and so Figure 6.8(2) is transformable to a figure $\mathcal{D}$, where $\mathcal{D}$ is obtained from Figure 6.8(2) by extending nation h to occupy the two holes touched by h. By figure $\mathcal{D}$, if there is a $w \in \{a, e\}$ such that edge $\{w, f\}$ is marked in G, then w is a correct crust of C. Similarly, if there is a $w \in \{a, e\}$ such that edge $\{w, h\}$ is marked in G, then the unique nation in $\{a, e\} - \{w\}$ is a correct crust of C. So, we may assume that none of the edges $\{a, f\}$, $\{e, f\}$, $\{a, h\}$ and $\{e, h\}$ are marked in G. Let G' be the marked graph obtained from $G - \{f, h\}$ by marking the following edges: $\{b, c\}$, $\{c, d\}$, $\{a, e\}$, $\{a, d\}$, $\{e, d\}$. By figure $\mathcal{D}$, we can obtain a well-formed atlas $\mathcal{M}'$ of G' from $\mathcal{M}$ by extending nation e^1 to occupy f and h. On the other hand, as in Case 1.1, we can prove that every well-formed atlas of G' can be used to construct a well-formed atlas of G.

7 Removing Maximal 4-Cliques

Throughout this section, we assume that G does not have an MC_5. We further assume that G has an MC_4; our goal of this section is to show how to remove MC_4's from G. The idea behind the removal of an MC_4 C from G is to try to find and remove a correct 4-pizza via constructing an extensible layout of C. After the removal of a correct 4-pizza, the resulting G may be not 4-connected and may have a separating 4-cycle, edge, triple, quadruple, or triangle. To restore Assumption 6, the algorithm reapplies the reductions in Sections 4 and 5 to the resulting G.

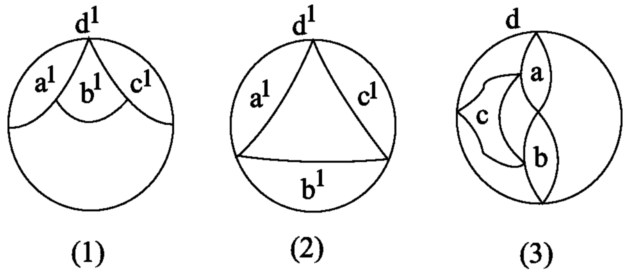

Figure 7.1: Possible displays of MC_4 $\{a, b, c, d\}$.

Suppose $C = \{a, b, c, d\}$ is an MC_4 of G; using Corollary 5.16 and $|V| > 8$, we find that only Figure 7.1(1), (2) or (3) can possibly display $\mathcal{M}_C$. Note these are a pizza, a pizza-with-crust, and a rice-ball, respectively.

7.1 Finding Rice-Balls

Let $C = \{a, b, c, d\}$ be an MC_4 of G. For a subset W of C, let $\mathcal{E}[W]$ be the set of unmarked edges $\{u, v\} \in E$ such that $u \notin W$, $v \notin W$, and some MC_4 of G consists of u, v, and two vertices in W.

Let $G' = G - C - \mathcal{E}[C]$. A *3-subset* of C is a subset S of C with $|S| = 3$. For each 3-subset S of C, let $V_S = \cup_K V(K)$, where K ranges over all connected components K of G' with $C \cap N_G(V(K)) = S$.

Lemma 7.1 *Figure 7.1(3) displays $\mathcal{M}|_C$ if and only if the following statements hold:*

1. *No two vertices in C are connected by a marked edge in G.*
2. *$V_{\{a,b,c\}}$, $V_{\{a,b,d\}}$, $V_{\{a,c,d\}}$, and $V_{\{b,c,d\}}$ each are nonempty and induce a connected component of G', and they together form a partition of $V - C$.*
3. *For every pair of two distinct 3-subsets S and T of C, $|V_S \cap N_G(V_T)| = 1$, $|V_T \cap N_G(V_S)| = 1$, and $(S \cap T) \cup (V_S \cap N_G(V_T)) \cup (V_T \cap N_G(V_S))$ is an MC_4 of G.*

Proof: For the "only if" direction, suppose that Figure 7.1(3) displays $\mathcal{M}|_C$. Then, $\mathcal{M}|_C$ has four holes, and each hole is touched by exactly three nations of C. For each 3-subset S of C, let $\mathcal{H}_S$ be the hole touched by the nations of S, and let Z_S be the nations of $V - C$ that occupy $\mathcal{H}_S$ in atlas $\mathcal{M}$. We want to prove that for each 3-subset S of C, $Z_S = V_S$. To this end, first observe that for each connected component K of G', there is a 3-subset S of C with $V(K) \subseteq Z_S$. This observation follows from Figure 7.1(3) immediately. Consequently, $C \cap N_G(V(K)) \subseteq S$; we claim they are equal. Towards a contradiction, assume that G' has a connected component K with $|C \cap N_G(V(K))| \leq 2$. Let $W = C \cap N_G(V(K))$. If $|W| \leq 1$, then K would be a connected component of $G - W$, a contradiction. If $|W| = 2$, then K is a connected component of $G - W - \mathcal{E}[W]$, and the vertices of W define a separating edge, a contradiction. So, the claim holds. By this claim, the above observation and Figure 7.1(3), we have $Z_S = V_S$ for each 3-subset S of C. In turn, by Figure 7.1(3), Statements 1 through 3 hold.

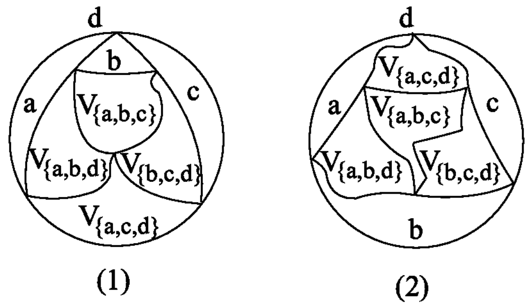

Figure 7.2: Possible atlases of G.

For the "if" direction, suppose that Statements 1 through 3 hold. We first prove that Figure 7.1(1) does not display $\mathcal{M}|_C$. Towards a contradiction, assume that Figure 7.1(1) displays $\mathcal{M}|_C$. We may assume that $\langle a^1, b^1, c^1, d^1\rangle = \langle a, b, c, d\rangle$ in Figure 7.1(1); the other cases are similar. Let S be a 3-subset of C. We claim that there is no point in $\mathcal{M}$ at which two nations u and v of V_S together with two nations x and y of $V - V_S$ meet cyclically in the order u, x, v, y. This claim holds; otherwise, $\{x, y\} \not\subseteq C$ by Figure 7.1(1), so x or y belongs to V_T for some 3-subset T of C other than S, and in turn $\{u, v\}$ would be a subset of $V_S \cap N_G(V_T)$, a contradiction against Statement 3 in the lemma. By this claim and Statement 2, the nations of V_S form a disc homeomorph in $\mathcal{M}$. Thus, by Statements 2 and 3, Figure 7.2(1) displays $\mathcal{M}$. By this figure, there is a 4-point p in $\mathcal{M}$ such that for each 3-subset S of C, exactly one nation $v_S \in V_S$ touches p. Since the nations v_S meet at p but no two of them belong to the same connected component of G', $C \cap N_G(v_S) = S$. In turn, by Figure 7.2(1) and the 4-connectedness of G, each V_S would equal $\{v_S\}$, contradicting Assumption 3. Therefore, Figure 7.1(1) does not display $\mathcal{M}|_C$.

We next prove that Figure 7.1(2) does not display $\mathcal{M}|_C$. Towards a contradiction, assume that Figure 7.1(2) displays $\mathcal{M}|_C$. As in the last paragraph, we may assume $\langle a^1, b^1, c^1, d^1\rangle = \langle a, b, c, d\rangle$, and we claim that the nations of each V_S form a disc homeomorph in $\mathcal{M}$. Thus, by Statements 2 and 3, Figure 7.2(2) displays $\mathcal{M}$. By this figure, there is a 4-point p in $\mathcal{M}$ at which nation a, some $u \in V_{\{a,b,c\}}$, some $v \in V_{\{a,b,d\}}$, and some $w \in V_{\{a,c,d\}}$ meet. Since u, v and w meet at p but no two of them belong to the same connected component of G', $C \cap N_G(u) = \{a, b, c\}$, $C \cap N_G(v) = \{a, b, d\}$, and $C \cap N_G(w) = \{a, c, d\}$. In turn, by Figure 7.2(2), $\{u\} = V_{\{a,b,c\}}$. Moreover, nations v, a, b, d meet at a point in $\mathcal{M}$, and nations w, a, c, d meet at a point in $\mathcal{M}$. Thus, $V_{\{a,b,d\}} = \{v\}$ or else $\langle u, b, v\rangle$ would be a separating triple of G, a contradiction. Similarly, $V_{\{a,c,d\}} = \{w\}$. In a similar way, we can also prove that $|V_{\{b,c,d\}}| = 1$. In summary, we have $|V| = 8$, a contradiction. Therefore, Figure 7.1(2) does not display $\mathcal{M}|_C$.

Since both Figure 7.1(1) and (2) do not display $\mathcal{M}|_C$, only Figure 7.1(3) can display $\mathcal{M}|_C$. This completes the proof. □

Since it is easy to check whether Statements 1 through 3 hold, we can easily decide whether C has an extensible "rice-ball" layout. Once we know that C has an extensible "rice-ball" layout, then by Figure 7.1(3) and Statement 2, we can easily find and then remove six correct 4-pizzas from G. By examining all the MC_4's in G, our algorithm can either find one that is a rice-ball, and thus make progress; or else it can establish that none of the MC_4's is a rice-ball.

7.2 Distinguishing Pizzas and non-Pizzas

By the previous discussion, we now suppose that our algorithm reaches a point where none of the MC_4's has a rice-ball layout. Then all the remaining MC_4's are either pizzas or pizza-with-crusts. Specifically, we have:

Corollary 7.2 *For every MC_4 C of G, and for every well-formed atlas $\mathcal{M}$ of G, either Figure 7.1(1) or (2) displays $\mathcal{M}_C$.*

Let $C = \{a, b, c, d\}$ be an MC_4 of G. Our goal in this section is to give a linear time decision procedure to decide which of Figure 7.1(1) and (2) displays $\mathcal{M}|_C$. Moreover, the procedure always chooses 7.1(2) when both are possible. Whenever Figure 7.1(2) displays $\mathcal{M}|_C$, we will have identified d^1 and therefore we immediately make progress by removing three correct 4-pizzas. When Figure 7.1(1) (the pizza) displays $\mathcal{M}|_C$, we do nothing with this MC_4 C and proceed to consider other MC_4's; this may eventually lead to a situation where all MC_4's in G have to be pizzas, as considered in Section 7.3.

If Figure 7.1(2) displays $\mathcal{M}|_C$, then there is no $e \in V - C$ with $\{a^1, b^1, c^1\} \subseteq N_G(e)$ or else V would equal $\{a, b, c, d, e\}$ by Corollary 7.2, a contradiction. Also, for Figure 7.1(2) to possibly display $\mathcal{M}|_C$, we must have (i) C is 3-sharing with exactly three MC_4 C_1, C_2 and C_3 of G and (ii) the unique nation of $C_1 \cap C_2 \cap C_3$ is adjacent to no nation of $V - (C \cup C_1 \cup C_2 \cup C_3)$ in G. We assume that $C_1 \cap C_2 \cap C_2 = \{d\}$, $C_1 = \{a, b, d, e\}$, $C_2 = \{a, c, d, f\}$ and $C_3 = \{b, c, d, g\}$; the other cases are similar. Let $U = \{a, b, \ldots, g\}$. By symmetry of the pizza, we may assume that $d^1 = d$ in Figure 7.1(1).

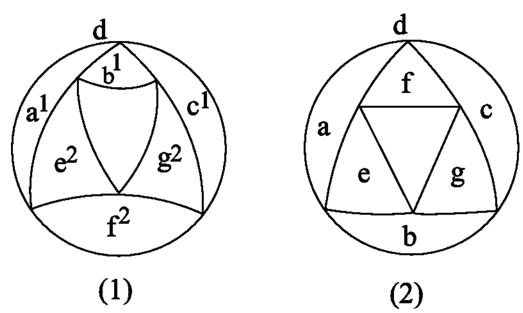

Figure 7.3: Possible displays of $G[U]$ when $\{e, f, g\}$ is a clique.

First, consider the case where $\{e, f, g\}$ is a clique in G. In this case, by Corollary 7.2, Figure 7.3(1) (or (2), respectively) displays $\mathcal{M}|_U$ if and only if Figure 7.1(1) (respectively, (2)) displays $\mathcal{M}|_C$. To distinguish the two figures, we check whether there is a nation $h \in V - U$ with $\{e, f, g\} \subseteq N_G(h)$. If no such h, then Figure 7.3(1) does not display $\mathcal{M}|_U$. If such h exists, then Figure 7.3(2) does not display $\mathcal{M}|_U$ because otherwise, V would equal $\{a, b, \ldots, h\}$ according to Corollary 7.2. In summary, when $\{e, f, g\}$ is a clique of G, we know which of Figure 7.1(1) and (2) displays $\mathcal{M}|_C$.

So, in the sequel, we assume that $\{e, f, g\}$ is not a clique of G. In case Figure 7.1(1) displays $\mathcal{M}|_C$, a simple inspection shows that one nation in $\{e, f, g\}$ (the one adjacent to a^1, c^1, d) is adjacent to the other two. So we assume that only one edge is missing among $\{e, f, g\}$, for otherwise Figure 7.1(2) must display $\mathcal{M}|_C$. We suppose the absent edge is $\{e, g\}$; the other cases are similar. Note that $\{a, d, e, f\}$ and $\{c, d, f, g\}$ are MC_4's in G.

Moreover, by Corollary 7.2, Figure 7.1(1) (or (2), respectively) displays $\mathcal{M}|_C$ if and only if Figure 7.4(1) (respectively, (2)) displays $\mathcal{M}|_U$. Figure 7.4(2) does not display $\mathcal{M}|_{U_i}$ if $\{d, f\}$ is a marked edge. Also, if $\{d, b\}$ is a marked edge, then Figure 7.5(1) does not display $\mathcal{M}|_{U_i}$ and so Figure 7.5(2) displays $\mathcal{M}|_{U_i}$. Thus, we may assume that neither $\{d, b\}$ nor $\{d, f\}$ is a marked edge.

To distinguish Figure 7.4(1) and (2), we do a case-analysis as follows:

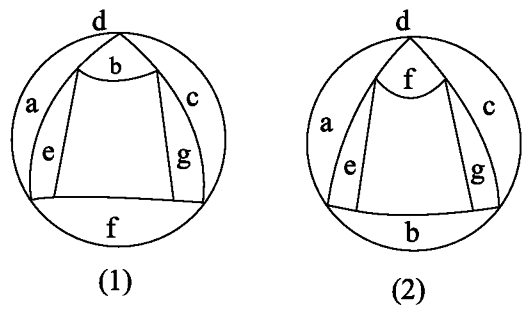

Figure 7.4: Possible layouts of $G[U]$ when $\{e, g\} \notin E$.

Case 1: There is no $h \in V - U$ with $\{a, b, e\} \subseteq N_G(h)$ or there is no $i \in V - U$ with $\{b, c, g\} \subseteq N_G(i)$. Then, Figure 7.4(1) does not display $\mathcal{M}|_U$. Whether h and i exist can be decided in $O(1)$ time, because $|N_G(a)| = |N_G(c)| = 6$ by Figure 7.4.

Case 2: There are $h \in V - U$ and $i \in V - U$ such that $\{a, b, e\} \subseteq N_G(h)$ and $\{b, c, g\} \subseteq N_G(i)$. Then, if $f \notin N_G(h)$ or $f \notin N_G(i)$, Figure 7.4(2) does not display $\mathcal{M}|_U$. So, we may assume that $f \in N_G(h)$ and $f \in N_G(i)$. Then, $h \neq i$ by Corollary 7.2, Figure 7.4(1), and (2). Let $U_i = U \cup \{h, i\}$.

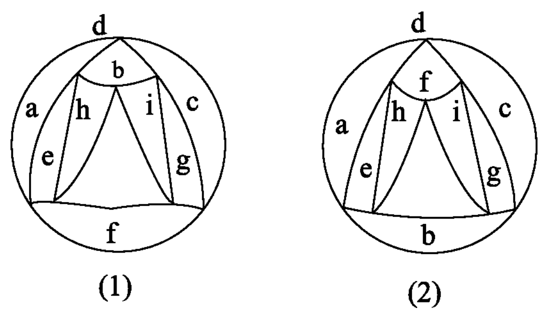

Figure 7.5: Possible layouts of $G[U_i]$ when $\{h, i\} \in E$.

Case 2.1: $\{h, i\} \in E$. If $N_G(f) \subseteq U_i$, then Figure 7.4(1) does not display $\mathcal{M}_U$ by Corollary 7.2. Similarly, if $N_G(b) \subseteq U_i$, then Figure 7.4(2) does not display $\mathcal{M}_U$. So, we may assume that neither $N_G(b) \subseteq U_i$ nor $N_G(f) \subseteq U_i$. Then, by Corollary 7.2 and Assumption 6 (G has no separating triple), Figure 7.5(1) (respectively, (2)) displays $\mathcal{M}|_{U_i}$ if and only if Figure 7.4(1) (respectively, (2)) displays $\mathcal{M}_U$. By Figure 7.5, $|N_G(e)| = |N_G(g)| = 6$; let j be the nation in $N_G(e) - U_i$ and k be the nation in $N_G(g) - U_i$. In case j or k is not adjacent to f in G, Figure 7.5(1) does not display $\mathcal{M}|_{U_i}$. Similarly, in case j or k is not adjacent to b in G, Figure 7.5(2) does not display $\mathcal{M}|_{U_i}$. So, we may further assume that j and k are adjacent to both f and b in G. Then, by Corollary 7.2, Figure 7.5(1) and

(2), we must have $j = k$ and $V = U_i \cup \{j\}$. Now, Figure 7.5(2) displays $\mathcal{M}|_{U_i}$ only if none of $\{a,b\}$, $\{b,c\}$, $\{b,h\}$, $\{b,i\}$, $\{e,f\}$, $\{f,g\}$, $\{f,j\}$ is a marked edge in G. On the other hand, if none of these edges is marked in G, then Figure 7.5(1) is transformable to (2) and hence the latter displays $\mathcal{M}|_{U_i}$.

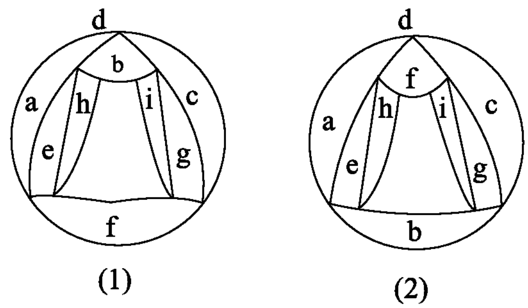

Figure 7.6: Possible layouts of $G[U_i]$ when $\{h,i\} \notin E$.

Case 2.2: $\{h,i\} \notin E$. Then, by Corollary 7.2, Figure 7.6(1) (respectively, (2)) displays $\mathcal{M}|_{U_i}$ if and only if Figure 7.4(1) (respectively, (2)) displays $\mathcal{M}_U$.

Case 2.2.1: There is no $j \in V - U_i$ with $\{e,f,h\} \subseteq N_G(j)$ or there is no $k \in V - U$ with $\{f,g,i\} \subseteq N_G(k)$. Then, Figure 7.6(1) does not display $\mathcal{M}|_{U_i}$. Whether j and k exist can be decided in $O(1)$ time, because $|N_G(e)| = |N_G(g)| = 6$ by Figure 7.6.

Case 2.2.2: There are $j \in V - U$ and $k \in V - U$ such that $\{e,f,h\} \subseteq N_G(j)$ and $\{f,g,i\} \subseteq N_G(k)$. Then, if $b \notin N_G(j)$ or $b \notin N_G(k)$, Figure 7.6(2) does not display $\mathcal{M}|_{U_i}$. So, we may assume that $b \in N_G(j)$ and $b \in N_G(k)$. Then, $j \neq k$ by Corollary 7.2, Figure 7.6(1), and (2). Let $U_k = U_i \cup \{j,k\}$.

Case 2.2.2.1: $\{j,k\} \in E$. Then, similarly to Case 2.1 above, we can distinguish which of Figure 7.6(1) and (2) displays $\mathcal{M}_{U_i}$.

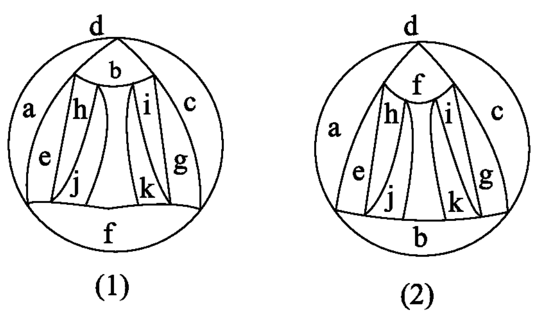

Figure 7.7: Possible layouts of $G[U_k]$ when $\{j,k\} \notin E$.

Case 2.2.2.2: $\{j,k\} \notin E$. Then, by Corollary 7.2, Figure 7.7(1) (respectively, (2)) displays $\mathcal{M}|_{U_k}$ if and only if Figure 7.6(1) (respectively, (2)) displays $\mathcal{M}_{U_i}$. In case at least one of the edges $\{a,b\}$, $\{e,f\}$, $\{c,b\}$, and $\{g,f\}$ is marked in G, Figure 7.7(2) does not display $\mathcal{M}|_{U_k}$. Moreover, in case at least one of the edges $\{a,f\}$, $\{e,b\}$, $\{c,f\}$, and $\{g,b\}$ is marked in G, Figure 7.7(1) does not display $\mathcal{M}|_{U_k}$. So, we may assume that no pair in $\{a,c,e,g\} \times \{b,f\}$ spans a marked edge of G.

Now, observe a resemblance between Figures 7.4 and 7.7. We want to iterate the above case-analysis to distinguish Figure 7.7(1) and (2). To this end, first observe that the above case-analysis is independent of nation d and edge $\{a, c\}$. Moreover, the case-analysis can be viewed as a procedure $CA(a, b, c, e, f, g)$ where the input parameters are nations of G related as in Figure 7.4(1) or (2) except for the possible absence of edge $\{a, c\}$. Thus, to distinguish Figure 7.7(1) and (2), it suffices to call $CA(h, b, i, j, f, k)$.

There can be a linear number of subsequent calls of procedure CA. Each call takes $O(1)$ time, so the overall time is linear.

7.3 Removing Pizzas

By the discussions in the last two subsections, we may assume that for every MC_4 $C = \{a, b, c, d\}$ of G, only Figure 7.1(1) displays $\mathcal{M}|_C$. That is, the four nations of every MC_4 of G meet at a point in $\mathcal{M}$.

Fix an MC_4 $C = \{a, b, c, d\}$ of G. C is 3-sharing with no MC_4 C' of G because otherwise, C' would have a non-pizza layout. By Figure 7.1(1), there are distinct nations e, f, g and h in $V - C$ such that $C \cap N_G(e) = \{a^1, b^1\}$, $C \cap N_G(f) = \{b^1, c^1\}$, $C \cap N_G(g) = \{c^1, d^1\}$ and $C \cap N_G(h) = \{d^1, a^1\}$, because $\mathcal{M}$ has no hole. On the other hand, the existence of the nations e, f, g and h ensures that the nations of C have to meet at a point in $\mathcal{M}$ cyclically in the order a^1, b^1, c^1, d^1. Thus, by finding out nations e, f, g and h, we can find and remove a correct 4-pizza from G.

By this method we may identify a correct 4-pizza for every MC_4 in G. Since these 4-pizzas all exist in every well-formed atlas of G, we may remove them all in one step by the remarks after Lemma 4.5.

8 Time Analysis

Let n and m be the number of vertices and edges in the input graph G, respectively. Suppose this is not a base case; that is, $n \geq 9$ and G has a 4-clique. Then we will show that the algorithm can always make progress in $O(n^2)$ time. In each case, the time needed to produce the subproblems from G dominates the time needed to recover a solution from the subproblem solutions, so we ignore the latter.

By Corollary 2.5 (with $k = 4$) G has $m = O(n)$ edges and arboricity $\alpha(G) = O(1)$, so we can list its $O(n)$ maximal cliques in linear time [7]. From the listed MC_4's, we can precompute the sets $\mathcal{E}[a, b]$ for all unmarked edges $\{a, b\}$, again in linear time.

We claim that testing the existence of a separating triangle takes $O(n^2)$ time. Since G has $O(n)$ maximal cliques and no 7-clique, it has $O(n)$ 3-cliques and these can be found in linear time. For each 3-clique C, it takes $O(n)$ time to test whether some (ordered) list of the vertices in C is a separating triangle. So, the claim holds. A similar analysis applies for finding a 3-cut (by Lemma 4.3(1)), a separating edge, or a separating triple.

In order to detect separating quadruples, we use an algorithm of Chiba and Nishizeki [7] which implicitly lists all 4-cycles of G in $O(m \cdot \alpha(G)) = O(n)$ time. The algorithm produces a list of triples (u_i, v_i, S_i) with the following properties:

1. u_i and v_i are non-adjacent vertices of G.

2. S_i is a set of vertices adjacent to both u_i and v_i.

3. Every induced 4-cycle in G occurs as $\langle u_i, x, v_i, y\rangle$ for some choice of i and $x, y \in S_i$.

In particular, the sum of all $|S_i|$ is $O(n)$.

We claim that testing the existence of a separating quadruple takes $O(n^2)$ time. It suffices to show the following: for each triple (u_i, v_i, S_i), we can test whether there is a separating quadruple $\langle u_i, x, v_i, y\rangle$ or $\langle v_i, x, u_i, y\rangle$ (with $x, y \in S_i$) in time $O(|S_i|n)$. By similarity, it suffices to show how to find those quadruples starting with u_i.

For x in S_i, let $G^x = G - \{u_i, v_i, x\} - \mathcal{E}[u_i, x]$. In linear time we may compute G^x and identify the set S^x of all cut vertices in G^x. Now there is a separating quadruple of the form $\langle u_i, x, v_i, y\rangle$ precisely if S^x contains some y which is in S but not adjacent to x. By repeating this for every $x \in S_i$, we have the required time bound.

A similar analysis applies for finding separating 4-cycles in $O(n^2)$ time.

The case analysis for eliminating an MC_5 in Section 6 may be executed in linear time. In particular, we may identify an MC_5 4-sharing with two other MC_5's in $O(n)$ time as follows. First, for every MC_5 C_i and for every $S \subseteq C_i$ with $|S| = 4$, create a pair (S, i). Next, bucket-sort all the pairs, and use the result to count the number of 4-sharing MC_5's with each C_i.

When the graph has no MC_5 but still has some MC_4's, we make progress in at most $O(n^2)$ time as follows. First, we list the $O(n)$ MC_4's in some arbitrary order. For each one, we test the conditions of Lemma 7.1 in $O(n)$ time; if we find such an MC_4, then we remove the identified 4-pizzas and we are done. Otherwise, we go through the list again, this time applying the linear time decision procedure of Section 7.2; if we determine that some MC_4 is a non-pizza, then we remove the identified 4-pizzas and we are done again. Otherwise, we have established that all the MC_4's are pizzas, and so we can remove a 4-pizza for each MC_4 by the method in Section 7.3.

Finally, if the algorithm reaches a base case, our graph G either has at most 8 vertices, or no 4-clique. In the former case we solve the problem exhaustively in $O(1)$ time. Otherwise, G should be planar; we finish in linear time [10], as described in Section 3.3.

Let $N = n + m$ be the *size* of our input graph, and let $T(N)$ be the maximum running time of the algorithm on any input of size N. We claim that there is a constant c such that $T(N) \leq cN^3$. The claim is clearly true for the base cases, as argued above. In all other cases, the algorithm makes progress in $c_1 N^2$ time for some constant c_1. That is, the algorithm produces one or more smaller marked graphs whose total size is larger than that of G by a constant c_2; the problem for G is reduced to solving the problem for each of these smaller instances. More precisely, there are integers $n_1, \ldots, n_\ell \in \{1, \ldots, N-1\}$ such that $\sum_{i=1}^{\ell} n_i \leq N + c_2$ and $T(N) \leq \sum_{i=1}^{\ell} T(n_i) + c_1 N^2$. We prove our claim by induction. For small N ($N < c_2^2$), our claim is true simply by choosing c large enough. For larger N, we have $T(N) \leq \sum_{i=1}^{\ell} cn_i^3 + c_1 N^2$ by the inductive hypothesis. Note that $\sum_{i=1}^{\ell} cn_i^3$ is maximized when $\ell = 2$, $n_1 = N - 1$ and $n_2 = c_2 + 1$. Hence, by choosing c large enough ($c_1 + 2$ suffices), we have $T(N) \leq cN^3$.

9 Concluding Remarks

Our main algorithm is too complex. We would like to find a faster algorithm, with simpler arguments. Perhaps such a simplification is possible using some of Thorup's ideas.

Naturally, we are very interested in polynomial-time algorithms for recognizing (hole-free or not) k-map graphs with $k \geq 5$. In view of the complication of our algorithm for hole-free 4-map graphs, however, new insights seem to be needed in order to make progress in this direction.

A natural and interesting question in connection with map graphs is to ask whether $\lfloor 3k/2 \rfloor$ colors suffice to color a k-map graph where $k \geq 3$. Note that $\lfloor 3k/2 \rfloor$ is the maximum clique size in a k-map graph. In case $k = 3$, the answer is positive because of the famous Four Color Theorem. As Thorup observed [15], the answer is also positive for $k = 4$: the 4-map graphs are all 1-planar graphs, and 1-planar graphs are known to be 6-colorable [2]. However, the answer is unknown when $k \geq 5$.

Similarly, we are interested in tighter versions of the bounds in Section 2; in particular the first author [3] has improved the edge bound in Corollary 2.5 to $kn - 2k$.

The recognition problem of map graphs is just a special topological inference problem, where each pair of regions either touch or are disjoint. One more general problem is obtained by allowing the relation between certain pairs of regions to be left unspecified (i.e., each such pair may touch or not touch). We conjecture that this generalization is NP-complete. Another generalization is obtained by allowing a region to include another region as a subregion. We conjecture that this generalization is polynomial-time solvable. Note that the inclusion relations among the regions should induce a rooted forest. The special case of this generalization where no four leaf regions meet at a point and each non-leaf region is the union of its descendant regions, can be solved by a nontrivial $O(n \log n)$-time algorithm [6]. In the real world, a non-leaf region is usually not a closed disc homeomorph; this more general problem is addressed in [5].

References

[1] J. F. Allen. Maintaining knowledge about temporal intervals. *Communications of the ACM*, 26(11):832–843, 1983.

[2] O. V. Borodin. A new proof of the 6 color theorem. *J. Graph Theory*, 19(4):507-522, 1985.

[3] Z.-Z. Chen. Approximation algorithms for independent sets in map graphs. Manuscript, 1999.

[4] Z.-Z. Chen, M. Grigni, and C. Papadimitriou. Planar map graphs. *Proc. 30th Ann. ACM Symp. Theory of Computing (STOC)*, 514–523, 1998.

[5] Z.-Z. Chen, X. He, and M.-Y. Kao. Nonplanar topological inference and political-map graphs. *Proc. 10th Ann. ACM-SIAM Symp. Discrete Algorithms (SODA)*, 195–204, 1999.

[6] Z.-Z. Chen and X. He. Hierarchical topological inference on planar disc maps. Manuscript, 1999.

[7] N. Chiba and T. Nishizeki. Arboricity and subgraph listing algorithms. *SIAM J. Computing*, 14(1):210–223, 1985.

[8] G. Ehrlich, S. Even, and R. E. Tarjan. Intersection graphs of curves in the plane. *J. Combinatorial Theory Ser. B*, 21(1):8–20, 1976.

[9] M. Grigni, D. Papadias, and C. H. Papadimitriou. Topological inference. *Proc. 14th Intl. Joint Conf. on Artificial Intelligence (IJCAI)*, 901–906, 1995.

[10] J. Hopcroft and R. E. Tarjan. Efficient planarity testing. *J. ACM*, 21(4):549–568, 1974.

[11] J. Kratochvíl. String graphs II: Recognizing string graphs is NP-hard. *J. Combinatorial Theory, Ser. B*, 52(1):67–78, 1991.

[12] J. Kratochvíl and J. Matoušek. String graphs requiring exponential representations. *J. Combinatorial Theory, Ser. B*, 53(1):1–4, 1991.

[13] C. St. J. A. Nash-Williams. Edge-disjoint spanning trees of finite graphs. *J. London Math. Soc.*, 36:445–450, 1961.

[14] J. Renz. A canonical model of the region connection calculus. In *Principles of Knowledge Representation and Reasoning: Proc. of the 6th Intl. Conf.*, 330–341, June 1998.

[15] M. Thorup. Map graphs in polynomial time. *Proc. 39th Ann. IEEE Symp. Foundations of Computing (FOCS)*, 396–405, 1998.