Nesting points in the sphere

Dan Archdeacon
Dept. of Computer Science
University of Vermont
Burlington, VT, USA 05405
e-mail: dan.archdeacon@uvm.edu

Feliu Sagols
CINVESTAV-IPN
Av. IPN 2508, Col Sn. Pedro Zacatenco
México, D.F. 07000
e-mail: fsagols@cs.cinvestav.mx

July 26, 1999 (revised June 14, 2000)

Abstract

Let G be a graph embedded in the sphere. A k-*nest* of a point x not in G is a collection $C_1, \ldots, C_k$ of disjoint cycles such that for each C_i, the side containing x also contains C_j for each $j < i$. An embedded graph is k-*nested* if each point not on the graph has a k-nest. In this paper we examine k-nested maps. We find the minor-minimal k-nested maps small values of k. In particular, we find the obstructions (under the minor order) for the class of planar maps with the property that one face's boundary meets all other face boundaries.

1 Introduction

Let G be a *spherical graph*, a graph drawn without crossings on the 2-dimensional sphere S. Our interest is in separating some points in the sphere from others using simple cycles in G. Of course, this is not always possible: two points in the same component, or same *face*, of $S-G$ are not separated. On the other hand, in some cases two points can be separated by many "nested" cycles of G. Let's make this more precise.

A *k-nest* of a point x in $S-G$ is a collection of simple cycles $C_1,\ldots,C_k$ of G such that for each i, the side of C_i that contains x also contains C_j for every $j<i$. Observe that if x is k-nested, then there exists a point y such that every xy-path in the sphere intersects G in at least k points; just pick y on the side of C_k that does not contain x. We say that the embedded graph G is *k-nested* if every point x in $S-G$ has a k-nest.

The definition above concerns only nesting of points in $S-G$. Call an embedded graph k^+-nested if every point of S is k-nested, including those points in G. This is a slightly stronger condition, as it is easy to see the following.

Lemma 1.1 *Every k^+-nested graph is k-nested, and every k-nested graph is $(k-1)^+$-nested.* ∎

A graph H is a *minor* of G if H can be formed from G by a sequence of edge contractions, edge deletions, and deletion of isolated vertices. The relation "*H is a minor of G*" forms a partial ordering on all spherical graphs.

Lemma 1.2 *If G has a minor H that is k-nested, then G is also k-nested. The same statement holds for k^+-nested graphs.*

Proof: The disjoint cycles $C_1,\ldots C_k$ for each point of the embedded H also serve as disjoint cycles of G. ∎

It is natural to consider the smallest graphs that are k-nested, or k^+-nested, where smallest refers to the minor order. In this paper we look for these graphs. The small cases are covered in Section 2, and the case of 2-nested graphs is covered in Section 3. We close with some concluding remarks in Section 4. These include an interesting rephrasing of Theorem 3.1 as a variation of outerplanar graphs. On with the proofs!

2 The small cases

In this section we consider the minimal k- and k^+-nested graphs for very small values of k.

Lemma 2.1 *The only minor-minimal* 0*-nested graph or* 0^+*-nested graph is* K_1*. The only minor-minimal* 1*-nested graph is a single loop on a vertex.*

Proof: The statement for 0- and 0^+-nested graphs is true because there is no restriction on the embedded graph. Any 1-nested graph must contain a cycle, and hence contains a single loop as a minor. ■

Having warmed up on the easier cases, we now show a preliminary lemma and then slightly harder proposition.

Lemma 2.2 *Let G be a minor-minimal k- or k^+-nested spherical graph. Then 1) G does not contain a cut edge, 2) G does not contain a degree two vertex, and 3) G does not contain three edges all in parallel. Moreover, if G is minor-minimal k-nested, then 4) G does not contain a degree 3 vertex incident with two parallel edges, and 5) G does not contain a degree 4 vertex incident with two pairs of parallel edges.*

Proof: The first statement holds because if e is a cut edge, the cycles of G correspond precisely to the cycles of $G - e$. Hence any nesting set for G is also a nesting set for $G - e$, contradicting the minimality of G. The second statement holds for the same reason when we consider G/e where e is incident with the vertex of degree two. For the third statement, let e be one of three pairwise parallel edges. Again, any nesting set of G corresponds to a nesting set of $G - e$, contradicting the hypothesized minimality of G.

For the fourth statement, if e is one of the two parallel edges incident with a degree three vertex, then any set of nesting cycles in G for a point not in G corresponds to a set of nesting cycles in the contracted G/e. (The vertex in G/e corresponding to e may not have the necessary nesting cycles, so the statement is false for k^+-nesting.) Likewise, for the fifth statement, any one of the four incident edges can be contracted and the graph remains k-nested, but not necessarily k^+-nested. ■

Proposition 2.1 *The only minor-minimal* 1^+*-nested graphs are those shown in Figure 1.*

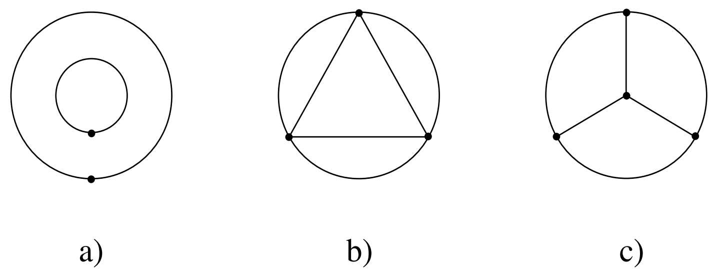

Figure 1: The minor-minimal 1^+-nested maps.

Proof: Let x be a vertex of G. Because x is nested, there exists a cycle C of G that does not contain x. If x is on a cycle that is disjoint from C, then we have a submap isomorphic to that on the left of Figure 1.

By Lemma 2.2 the degree of x in G is at least three. If there exist paths from x to three different vertices of C, then G contains the K_4 minor shown on the right of Figure 1. Hence the component of $G - C$ that contains x can attach to at most two points on C; call these points y and z. Observe that this component must attach to both y and z, or else we have either disjoint cycles or three pairwise parallel edges.

Now, the degree of x is at least three, so there exist two distinct paths from x to say y and one from x to z. Every point on the sphere except y is now disjoint from a cycle. Let C' be the cycle disjoint from y. Because there are no disjoint cycles, C' must intersect both the digon on x, y and the digon on z, y. Add a subpath of C' between these two digons. The resulting graph contains the graph shown in the middle of Figure 1 as a minor. ■

Having finished the "easy" cases, in the next section we will look at a harder case.

3 Minimal 2-nested graphs

In this section we want to find all minor-minimal 2-nested spherical embedded graphs. They are shown in Figure 2. Some of these spherical graphs shown have variants (1-flips) that are also minor-minimal 2-nested spherical graphs; we'll make this more precise shortly.

The following lemma is useful when considering 2-nested graphs. Its proof follows from the fact that the boundary walk of any face of an embedded graph that is not a tree contains a cycle.

Lemma 3.1 *A graph G embedded in the sphere is 2-nested if and only if for every face f, there is a face g such that their boundary walks in G are disjoint.* ■

Let e be a loop incident with a vertex v. Consider an embedding of $G-e$ in the sphere. There are many different ways to extend the embedding to include e. For example, the edge e could be added in a small neighborhood of v in any incident face. If two embeddings of the same graph differ only in where a loop attaches at a vertex, we say that one is a *1-flip* of the other. The following is not difficult.

Lemma 3.2 *Let G be a minor-minimal 2-nested graph in the sphere. Then any 1-flip of G is also minor-minimal 2-nested.*

By Lemma 3.2 we need only consider minor-minimal 2-nested graphs up to 1-flips. If we consider 1-flips as different maps, then the number of minimal 2-nested maps increases. For example, G_2 has a 1-flip which is not isomorphic as a map.

We make even a finer distinction between maps. When considering when two maps on the sphere are isomorphic, it is sometimes convenient to consider maps on the oriented sphere and to consider only orientation-preserving homeomorphisms of the sphere to determine isomorphic maps. In this scenario, a map is not necessarily isomorphic to its mirror image (the same map with the opposite orientation). This can introduce some additional minor-minimal oriented maps.

We now give our main result.

Theorem 3.1 *There are exactly 9 minor-minimal 2-nested spherical graphs up to 1-flips. They are given in Figure 2. There are exactly 12 minor-minimal 2-nested spherical graphs up to map isomorphisms. There are exactly 16 minor-minimal 2-nested oriented spherical graphs up to oriented map isomorphisms.*

Proof: The last two statements follow from the first one with some tedious case checking. Also, it is tedious but straightforward for the reader to verify that each of the 9 graphs are 2-nested, but do not remain so after the deletion

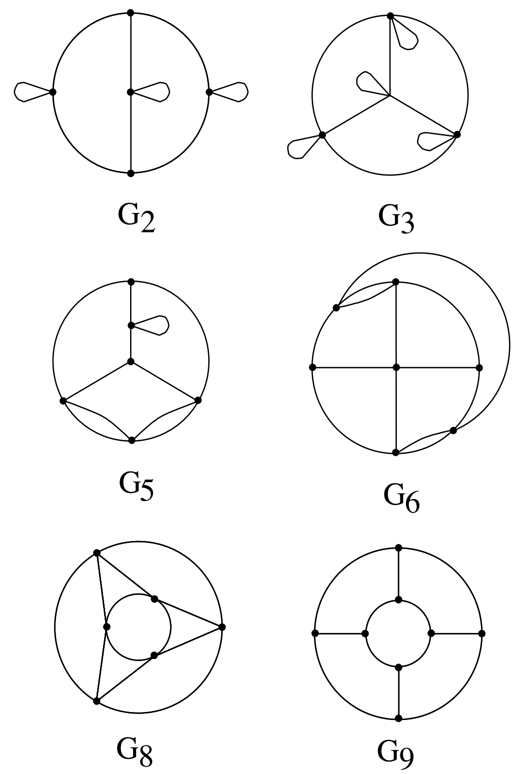

Figure 2: The minor-minimal 2-nested maps up to 1-flips.

or contraction of any edge. The remainder of the proof of the first statement follows from Propositions 3.1–3.6. ∎

We now give the proofs of the propositions. The names of the graphs come from their labels in Figure 2.

Proposition 3.1 *Let G be a minor-minimal 2-nested spherical map that is not connected. Then G is G_1.*

Proof: Let K be a component of G such that one side, which we'll call the inside, of K has no other component of G. Let x be a point on the outside of K very near an edge of K. This x is 2-nested. If both cycles are in K, then these two cycles together with a cycle from another component give a G_1 submap. If both cycles are outside of K, then these together with a cycle of K give a G_1 submap. ∎

Proposition 3.2 *Let G be a connected minor-minimal 2-nested spherical map that contains a loop. Then G is either G_2, G_3, or G_5.*

Proof: Let e be a loop on a vertex v. Find a 1-twist of e, if necessary, so that e bounds a face. Let x be a point in the other face with e in its boundary. Then x is 2-nested by cycles C_1 and C_2. If the loop e is disjoint from C_1 and C_2, then we have a G_1 subgraph. But e cannot intersect C_2 nontrivially, because it is separated from C_2 by C_1. Hence C_1 must contain v, and we have the subgraph shown on the left half of Figure 3. Following that figure, we shall refer to the face bounded by e as the inside, the side of C_1 containing C_2 as its inside, and the side of C_2 not containing C_1 as its inside.

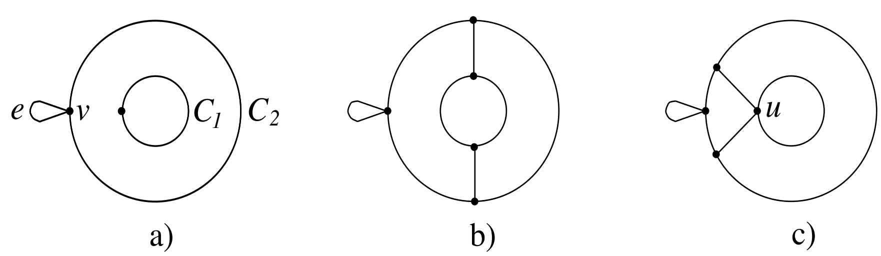

Figure 3: The graphs of Proposition 3.2.

We can assume that we pick C_1 and C_2 among all possible such subgraphs so as to enclose the maximum possible number of faces of G in the annulus bounded by these cycles. Observe that any point inside of C_2, outside of C_1, or inside the loop e is already 2-nested. The only points not yet 2-nested are in the annulus bounded by C_1 and C_2. Now $G - e$ is not 2-nested, so there exists a face f (necessarily in the annulus) whose boundary does not contain v but intersects every other face in the annulus. Hence either there is a vertex u on C_2 that separates C_1 from the rest of G, or there is a submap H isomorphic to either the middle or right of Figure 3.

If there is the cut-vertex u, then let g be the region in the annulus that contains both v and u. To nest the points in g there exists a cycle disjoint from v and from the boundary of g. This gives a G_1 minor, a contradiction.

The case shown on the right of Figure 3 is a minor of the case shown in the middle of that figure. As it happens, the same argument works for both cases. We will argue off of the rightmost figure and leave the extension to the other case for the reader. Let H be the subgraph of G shown in Figure 3.

Let g be the region of H that is in the annulus and contains v on the boundary. Points x in g are the only ones not two-nested. If there exists a face h disjoint from g, then there are three possibilities. First, if h is outside C_1, or if h is inside of C_1 but its boundary does not intersect C_2, then we have a G_1 minor. If h is inside of C_1 and its boundary does not intersect C_1, then we have either a G_1 or G_2 minor. Finally, if h intersects both C_1 and C_2, then we have a G_5 minor. We conclude that there is no face that is disjoint from the boundary of g, and in particular, we conclude that g cannot be a face of G.

Over all possible such subgraphs H, pick the one where g has the minimal possible number of faces of G. Then, avoiding a G_1 minor, there is no chord in g which is not incident with v. In fact, any chord must have as a minor a subgraph with chord uv. Let g_1 be the subregion of $G \cup uv$ "above" uv, and g_2 be the subregion "below". Now, g_1 must have a face whose boundary is disjoint from the boundary of g_1, but by the above intersects the boundary of g. Likewise g_2 has a face whose boundary is disjoint from g_2 but which intersects g. It follows that G has a G_3 minor, and the proposition is shown. ■

Proposition 3.3 *Let G be a connected minor-minimal 2-nested loopless spherical graph that contains a subgraph homeomorphic to the left half of Figure*

4. Then G is G_4.

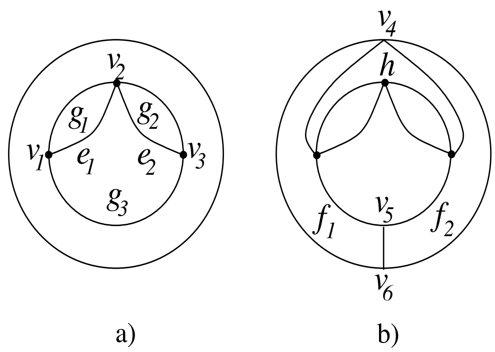

Figure 4: The graphs of Proposition 3.3.

Proof: We refer the reader to Figure 3.3 for information on how our subgraph H is labeled. We use that figure to define the inside and outside of the two circuits in the obvious manner. Note that any point outside of C_2, inside of C_1, or in either of the two 2-sided faces is already 2-nested. It remains only to find the additional parts of $G - H$ that make points in the annulus 2-nested.

The points in g_1 remain 2-nested in the graph $G - e_1$. Since there are points in $G - e_1$ that are not 2-nested, they must lie in a face f_2 whose boundary is disjoint from that of g_1, but which intersects the boundary of all other faces of G. Similarly, deleting e_2 gives a graph which is not 2-nested, so these points lie in a face f_1 whose boundary is disjoint from g_2 but which intersects the boundary of all other faces. These faces f_1 and f_2 are necessarily distinct and lie in the annulus bounded by C_1 and C_2, so G has as a minor the subgraph shown in the right half of Figure 4.

The only points that are not now 2-nested lie in the face labeled h. Hence there exists a face whose boundary does not intersect that of h. This face's boundary must intersect both cycles in the annulus, or else it can be contracted to a loop. We now have G_4 as a minor, as desired. ■

Proposition 3.4 *There is no G that is a minor-minimal 2-nested spherical graph that contains a topological K_4 disjoint from a cycle.*

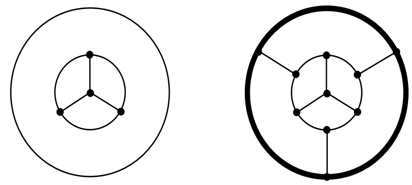

a) b)

Figure 5: The graphs of Proposition 3.4.

Proof: By way of contradiction suppose that G was such a graph with a subgraph H isomorphic to the left side of 5. Chose H so that the number of faces of G in the annulus between the two components is maximal. Call the three edges of the K_4 which are not on the boundary of the annulus *spokes*. If we contract any of the three spokes, the resulting graph is not 2-nested. Hence there must exist a face in the annulus incident with every other face except the K_4 face incident with the other end of the spoke. It follows that G contains a graph as in the right half of 5, where some of the edges in the outside cycle may be contracted. This graph has a cube minor by contracting the entire outside cycle to a point, a contradiction that G was chosen minor-minimal. ∎

A *theta-graph* is a subgraph homeomorphic to one with two vertices and three parallel edges joining them. It is called this because it resembles the Greek letter θ.

Proposition 3.5 *Let G be a minor-minimal 2-nested spherical graph that contains a theta-graph disjoint from a cycle, but is not covered under previous propositions. Then G is G_6 or G_7.*

Proof: We label the graph as depicted in Figure 6a and refer to the inside of C_1 and C_2 as the sides containing the remaining edge e. The only points that are already not 2-nested by this subgraph H are those in the annulus bounded by C_1 and C_2. Chose the copy of H in G so that the number of faces of G in this annulus is maximized.

Because G is minor minimal, G/e is not 2-nested. Hence there exist faces f_1 and f_2 in the annulus that are incident with e and at least one of them has a boundary incident with all other faces. If either one of the two faces is not incident with C_2, then we fall in to the case of Proposition 3.3 or Proposition

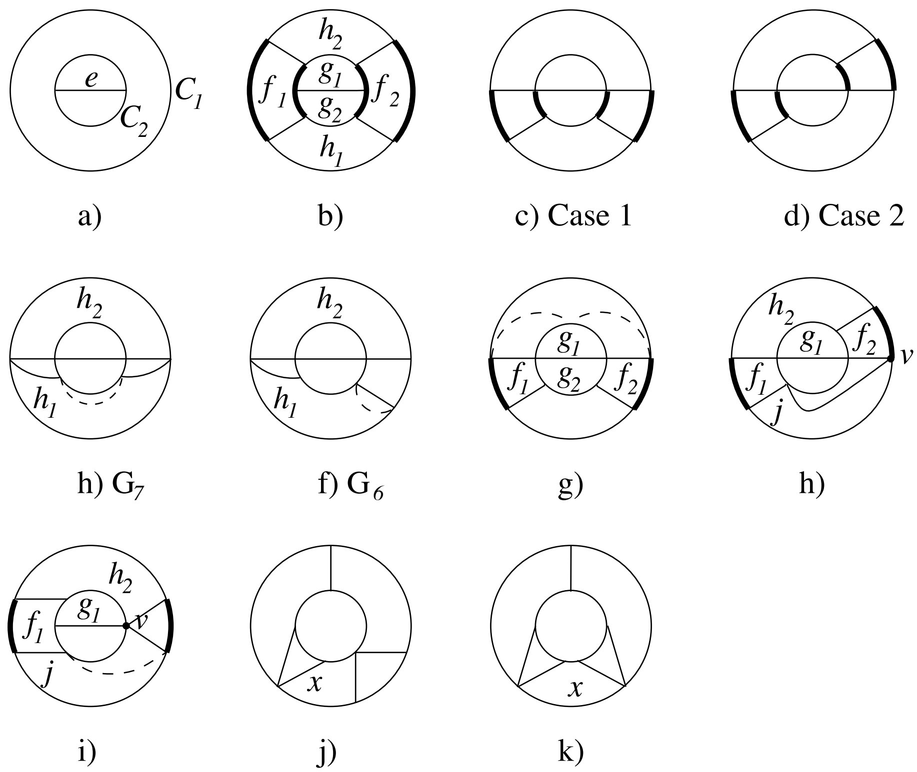

Figure 6: The graphs of Proposition 3.5.

3.4. Hence G must have a subgraph H as depicted in Figure 5b. The bold lines may be edges of H, or may be contracted to vertices of H. Label the faces f_3 and f_4 as shown.

If none of the bold edges are contracted, H has many cube minors. But the cube is our G_9, hence it cannot be contained in any other minor-minimal example. We break into cases.

Case 1: *G contains the graph of Figure 6c.* If we delete e, the resulting graph is not 2-nested. Hence there exists a face j either 1) its boundary does not intersect that of g_1, but intersects all other faces, or 2) its boundary does not intersect that of g_2, but intersects all other faces. These form two subcases.

Case 1.1: The face j of G must lie in the face h_1 of H, because no other face of H has the possibility of intersecting the boundary of all faces except g_1. Moreover, one or both of the bold edges on C_2 must be contracted. If both edges are contracted we get the graph of Figure 6e, without the dotted edge. The points in h_2 are not 2-nested. The only possibility for a disjoint face that meets all other faces has as a minor the graph of Figure 6e with the dotted edge. This is G_7 on our list. If only one of the bold edges on C_2 is contracted, then we get the graph of Figure 6f, without the dotted edge. Again, the only way to choose a second face disjoint from h_2 that avoids the previous case has as a minor the graph with the dotted edge of Figure 6f, which is G_6 on our list.

Case 1.2: The face j of G must lie in the face h_2 of H. Its boundary must intersect that of g_1, f_1, and f_2, and must not intersect that of g_2. The only possibility has as a minor the graph of Figure 6g. This graph has an octahedral minor formed by contracting the bold edge and deleting one of the two resulting parallel edges. The octahedron is G_8 on our list, contradicting that G was minor-minimal.

Case 2: *G contains the graph of Figure 6d.* If we delete e, the resulting graph is not 2-nested. Hence there exists a face j such that, without loss of generality, its boundary is disjoint from g_1 but intersects the boundary of all other faces. This gives the graph of Figure 6h as a minor. As before, the two bold edges may be contracted in this figure. We break into two subcases, depending on whether the bold edge in the boundary of f_2 is 1) contracted, or 2) not contracted.

Case 2.1: After the bold edge on f_2 is contracted, we examine why we cannot contract the edge between faces g_1 and f_2. There must exist a face incident with one end of this edge whose boundary intersects all other faces except one on the other end of the edge. One of these two faces cannot contain the vertex v. If the face without v is in not in h_2, then we get either the subgraph of Proposition 3.3 or of Proposition 3.4. Hence the face is in h_2. We now get an octahedral minor (G_8) as in Case 1.2.

Case 2.2: If the edge on f_2 is not contracted, then we contract the edge between f_2 and h_2 and the other solid edge and get an octahedral minor.

Case 3: *G contains the graph of Figure 6i (without the dotted line).* Now, $G - e$ is not 2-nested, so there exists a face j disjoint from (without loss of generality) g_1 and whose boundary intersects all other faces. The only possibility for j is the labeled face with the dotted line included where at least one of the two bold edges are contracted.

The graph formed from contracting the edge between f_1 and g_1 is not 2-nested, so there are two faces k_1 and k_2 incident with either end of that edge whose boundaries are disjoint. Moreover, one of these faces, say k_1, has a boundary that intersects all other face boundaries except k_2. This face k_1 must lie in h_2 and its boundary must contain the vertex v. The other face k_2 must lie in either g_1, g_2, or f and cannot contain v. All possibilities for k_2 give the cases of Propositions 3.3 or 3.4.

Case 4: *None of the previous cases.* Observe that if none of the bold edges are contracted, then G has two pentagons f_1 and f_2 with five pairwise non-adjacent edges joining them. The only ways to avoid a cube minor and not have the contractions of the first three cases are shown in Figure 6j and 6k. The only points not yet nested are in the region with x. Any way of nesting these two points gives a previous case.

This ends the casework and completes the proof of Proposition 3.5. ■

Proposition 3.6 *Let G be a connected minor-minimal 2-nested spherical graph without loops that does not contain a theta-graph disjoint from a cycle. Then G is G_8 or G_9.*

Proof: Pick an arbitrary point x and use the fact that it is 2-nested to find disjoint cycles C_1 and C_2 in G. Suppose that there is a vertex v of G that is not in $C_1 \cup C_2$. There cannot be paths from v to two distinct points of C_1, or else we get a theta-graph disjoint from a cycle. Similarly there cannot be paths from v to distinct points in C_2. There cannot be three distinct paths from v to C_1, or else again there is a theta-graph disjoint from a cycle. Hence either v is of degree two, degree 3 and incident with two parallel edges, or incident with two faces each bounded by digons. Each of these cases are impossible by Lemma 2.2. We conclude that all edges of $G - (C_1 \cup C_2)$ have one end in C_1 and the other in C_2.

We consider the faces in the annulus between C_1 and C_2. First, suppose that there is a face bounded by two edges in the annulus. Let f be another face with one of these two edges in its boundary. There exist a face g in the annulus disjoint from f. This face g is also disjoint from the digon, and hence is disjoint from a theta graph. So all faces are bounded by at least three edges.

Second, suppose that there exists a quadrilateral annulus face disjoint from another quadrilateral face. Then G contains the cube, G_9.

Third, suppose that there is a quadrilateral face in the annulus that is disjoint from a triangle. Then we have a subgraph H as shown in Figure 7a, except for the dotted lines. Now, the quadrilateral and the triangle are disjoint cycles, and by the argument above all remaining edges must have an end in each cycle. The only possible edges that do this and are not incident with the vertex labeled v are shown as dotted edges in Figure 7a. Both of these dotted edges must be included to 2-nest all points in the annulus on faces incident with v. But the resulting graph has an octahedral minor—G_8 on our list—and hence is not minor-minimal.

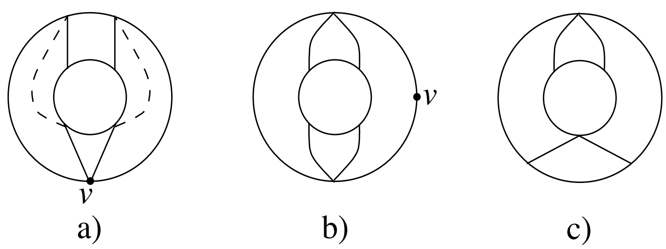

Figure 7: The graphs of Proposition 3.6.

¿From the first three cases we conclude that there is no quadrilateral face, in fact, that all faces are triangles.

Fourth, suppose that there exist two disjoint triangular faces in the annulus that both have an edge on C_1 (or symmetrically on C_2), giving the graph on Figure 7b. The outside face cannot be bounded by two parallel edges, since if so G is not 2-nested. Hence G has a vertex v as shown. Any edge from v to C_1 creates a theta-graph disjoint from a cycle, contradicting the hypothesis of this proposition.

We are down to the case that all faces are triangles, and all disjoint pairs of triangles in the annulus have one edge from C_1 and one edge from C_2. Starting with the picture in Figure 7c it is easy to see that G contains the octahedron G_8 as a minor. ■

4 Conclusion

We begin by rephrasing Theorem 3.1. Recall that by Lemma 3.1, a graph is 2-nested if and only if for every face f there is a face g such that the boundaries of f and g are disjoint. Hence graphs that are not 2-nested have the property that there is at least one face f whose boundary meets the boundary of all other faces. This is a variation on outerplanarity: the special face need not contain every vertex, but must contain at least one vertex from all other faces. Hence we have shown the following.

Corollary 4.1 *A spherical map G has a face whose boundary intersects all other face boundaries if and only if it does not contain one of the maps of Figure 2 or their 1-flips as a minor.*

We could have considered other orderings besides graph minors. For example, a graph H is a *topological subgraph* of G if H can be formed by deleting edges, deleting isolated vertices, and replacing the two edges in parallel on a degree two vertex with a single edge. The topological ordering is coarser than the minor ordering, hence there might be more topologically-minimal k-nested graphs than minor-minimal k-nested graphs. Any topologically minimal graph can be transformed into a minor-minimal graph by contracting edges. It would be a straightforward task to take the graphs of Figure 2 and construct the topologically-minimal 2-nested graphs, but in this paper we chose to focus on the minor order.

The relationship between graph minors and structural properties such as tree-width has been widely studied recently [3, 4]. For the appropriate definitions and a excellent survey we refer the reader to Reed [1]. The authors originally hoped for a relationship between having a high nesting number and having large tree-width. But a large nesting number might have small tree-width: a set of $2k-1$ nested loops is a k-nested graph with tree-width 2. However, these graphs have a special properties that leads to the next definition.

Let G be a spherical graph, and let x be a point of $S - G$. The *local nesting* number of x, $\nu(x)$ is the maximum k with nested cycles $C_1, \ldots, C_k$ around x. The *roundness* of G is the maximum ratio $\nu(x)/\nu(y)$ over all points x, y. A set of $2k-1$ nested loops has roundness close to 2. The roundness is bounded between 1 and 2.

Conjecture 4.1 *Let $\mathcal{G}$ be the family of graphs with roundness bounded above by $2-\epsilon$. Then there exists a function $f(n)$ such that for any $G \in \mathcal{G}$, a nesting number at least $f(n)$ implies a tree-width at least n.*

ACKNOWLEDGEMENTS: The authors thank Thomas Fowler and Barrett Walls for helpful discussions on this topic. We also thank the referee who found an error in the original statement of Lemma 2.2.

References

[1] B.A. Reed, Tree width and tangles: a new connectivity measure and some applications. Surveys in Combinatorics, 1997 (London) 87-162, *London Math. Soc. Lecture Note Ser.*, **241** Cambridge Univ. Press, Cambridge, 1997.

[2] G. Chartrand and F. Harary, Planar permutation graphs. *Ann. Inst. Henri Poincaré Sec. B* **3** (1967)433-438.

[3] N. Robertson and P.D. Seymour, Graph Minors X: Obstructions to tree-decomposition, *J. Combin. Th., Ser. B* **52**(1991)153-190.

[4] P.D. Seymour and R. Thomas, Graph searching and a min-max theorem for tree-width, *J. of Comb. Th., Ser. B* **58**(1993)22-33.