COMP599 Assignment 2 Solutions

1(a) Let h(k) be max(4,g(k,k)+1) for the function g of Lemma 4.2. Suppose that the desired result is false and consider a counterexample (G,H,X,v) chosen to minimize |V(H)| +|E(H)|. Let $(Y_1,Y_2,\ldots,Y_p)$ be a partition of X which is realizable in G but not G-v. Let $T_1,\ldots,T_p$ be a realization of this partition in G chosen so as to minimize the number of vertices in these trees. Clearly v is in some T_i wlog T_1.

If w is a vertex of G-X-v in no C_i and no T_i then deleting it contradicts the minimality of G. If e is an edge of G which is in no C_i and no T_j then deleting it contradicts the minimality of G.If e is an edge which is in some C_i and does not join two distinct trees of our realization then contracting it contradicts the minimality of G. It follows that the edges of the realization are precisely those edges which are not in some C_i and the realization is unique. It also follows that the set of cycles is unique. It also follows that every vertex is in some T_i. Finally, if e is an edge of G which is in some T_i and does not either join v to C_1, join C_i to C_j for distinct i and j, or join two vertices of $C_{h(k)}$ then contracting it contradicts the minimality of G. It follows from this that the vertex set of G consists of the union of the vertices of the cycles and v. It also follows that v has degree 2.

If l is a leaf of some T_i which is not in X then replacing T_i by T_i-l contradicts our choice of a realization. It follows that every vertex w not in X is joined to X by two paths of some T_i disjoint except at w. In particular, there are two paths P_1 and P_2 from v to X which intersect only at v. .

Note that no edge of G which is not in H has an endpoint in V(H)-$C_{h(k)}$. Thus, each C_i separates v from X. It follows that both P_1 and P_2 must pass through every C_i (before leaving the closed disc bounded by C_i).Thus G is 2-connected

Let H' be the planar graph obtained from H by contracting the component of H-$C_{h(k)-1}$ containing $C_{h(k)}$ into a single vertex b and deleting v and adding an edge between its two neighbours on C_1. We also modify the tree of our realization containing v by deleting v and adding an edge between v's neighbours. We note that b is adjacent only to vertices of $C_{h(k)-1}$ in H'.

Assume that H' is not 3-connected. Let Z be an arbitrary minimal cutset of size at most 2 in H'. Thus every vertex of Z has a neighbour in each component of H'-Z.

If b is in Z then Z intersects $C_{h(k)-1}$ in at most one vertex so $C_{h(k)-1}$-Z is connected and lies in one component of H'-Z. Thus, b contains a neighbour in the other component of H'-Z which is disjoint from $C_{h(k)-1}$, a contradiction.

So b is not in Z, and Z is a cutset in H such that some component U of H-Z completely contains $C_{h(k)}$. Thus, by the definition of H, every other component of H-Z is a component of G-Z. Let K be any other component of G-Z and let w be a vertex of K. Now K is disjoint from X, so K contains no leaf of any T_i. Since w is not in X, there are two paths of T_i from w to X, disjoint except at w, which must pass through Z. Thus there is a path Q of T_i whose interior is in K, joining the two vertices of Z. It follows that there is at most one choice for K, and hence G-Z has two components. Since there are no leaves of any T_j in K, we see that K intersects our realization precisely in the interior of Q. This implies that every vertex of K is on the interior of Q. It also implies that for any other path Q' joining the vertices of Z whose interior is in Q, replacing T_i by T_i-Q+Q' yields a realization of the partition $(Y_1,..,Y_p)$. Thus, by the

uniqueness of the partition, Q is induced. It follows that if a C_j intersects K then Q is a subpath of C_i. But now contracting any edge of Q which is not incident to v contradicts the minimality of G (it is in some T_i does not join disjoint cycles and does not touch v). It follows that Z consists of the neighbours of v, and V(K) is v. But now, Z is not a cutset of H' because v is not in H'.
It follows that H' is a a 3-connected graph which has a unique embedding. Thus so does the graph H* obtained by adding back v in the middle of the edge of H' between its neighbours via a subdivision. We can take obtain this embedding by either contracting from our our planar embedding of G or by contracting from the planar embedding of H which shows that v is h(k)-protected. Thus, just as in the latter embedding, in the planar embedding of G, C_1 to $C_{h(k)-1}$ are nested cycles separating v from $C_{h(k)}$.

Thus, since there are no vertices of G apart from v which do not lie in one of the nested cycle,
The boundary of every face incident to a vertex of X intersects $C_{h(k)}$. For each vertex w in X, we choose a face whose boundary contains w. Draw a small disc in that face which intersects G only at w. Make these discs small enough that they do not intersect. The desired result now follows from Lemma 4.2.

(b) We consider a minimal counterexample and minimal realization as in 1(a). We again derive that: every edge either is in a C_j and joins distinct trees of the realization or is in a T_i and joins distinct cycles, or joins v to C_1. This implies that the realization is unique, the set of protecting cycles is unique and that V(G) is V(H). We also again obtain that every vertex w not in X is joined to X by two paths of some T_i disjoint except at w. Applying this to v, we see that G is 2-connected..

We will show that in fact G is planar. To do so, we show that the Lipton-Tarjan decomposition splits B into triangles and one 3-connected graph which has no (3,3)-cut which, by Wagner's Theorem, completes the proof.

We begin with a detour. We ask if there are three vertex disjoint paths of G from C_1 to X. If not there is a cutset Z of size 2 separating C_1 from X. Since C_1 separates v from X, Z also separates v from X. We let A be the union of those components of G-Z which contain v or a vertex in C_1. Since X is disjoint from A, the restriction of our realization to the subgraph of G induced by A cup Z is a path Q between the vertices of Z. Since the realization is unique there is only one such path with interior in A. But there are two disjoint paths from C_1 to Z which are subpaths of P1 and P2 and hence two paths with endpoints in Z and interior in A (we can go either way around C_1), a contradiction. So the desired three paths exist. We note that they actually exist in G-v since C_1 separates v from X.

We next ask if there are four vertex disjoint paths of G from C_1 to X. If not there is a cutset Z of size 3 separating C_1 from X. Since C_1 separates v from X, Z also separates v from X. Since both of v's neighbours are on C_1 and not in X we have that v is not in Z (as otherwise Z-v would separate C_1 from X, contradicting the last paragraph). We let A be the union of those components of G-Z which contain v or a vertex in C_1. Since X is disjoint from A, the restriction of our realization to the subgraph of G induced by A cup Z is a forest F whose leaves are in Z. But the fact that there are three paths from X to C_1 in G-v imply that there are 3 paths from Z to C_1 with their interior in A-v. These along with the cycle C_1 show that there is a forest F' in A cup Z –v with the same leaves as F such that replacing T_1 with T_1-v yields a realization of $(Y_1,\ldots,Y_p)$ in G-v, a contradiction Hence the desired four paths do exist.

Now consider some cutset Z of size at most 3 in G. There is a component K_Z of G-Z completely containing some C_j. K_Z intersects all four of the disjoint paths of the last paragraph and completely

contains one of them (which is disjoint from Z). Thus it also completely contains C_j-Z for every cycle not intersecting Z in 2 vertices, and intersects every C_j. It will be useful to know this in what follows.

Suppose now that G has a cutset Z of size 2. We let K be some component of G-Z other than K_Z. If K is disjoint from every C_i then the minimality of G implies that it is v. If K intersects some C_i then, since it does not contain all of any C_j it must contain a path Q of C_i between the two vertices of Z, and no other vertices of any C_j. Now every edge of C_i joins distinct trees of our realization, applying this to edges of C_i from Z into K, we see that there is a T_j intersecting K which contains at most one vertex of Z. Since every vertex w of T_j-X is joined to X by two paths which are disjoint except at w, we see that X intersects K and thus i is h(k). We know that K_Z contains (i) all of every C_j with j less than h(k), (ii) a vertex of $C_{h(k)}$ and hence a path of $C_{h(k)}$ between the vertices of Z, and (iii) v (since Z is disjoint from C_1). Thus the vertex set of K is precisely the interior of Q. We note that, since K intersects H precisely in the interior of Q, and there are no edges from K to H-$C_{h(k)}$ replacing Q by any path between the vertices of Z through K would yield a new set of protecting cycles. Since the set of protecting cycles is unique, Q is the unique such path and it is induced.

We have shown that every 2-cut of G has 2 components one of which is a path, all of whose vertices have degree 2 in G. . Hence, G is the subdivision of a 3-connected graph G'.

Suppose now that G and hence G' is not planar. G' is clearly not L, as G' contains too many nested circuits (L does not contain 3 disjoint circuits). Thus, by Wagner's theorem G' has a (3,3)-cut Z which is also a (3,3) cut of G. Let K_1=K_Z ,K_2, and K_l be the l> 2 components of G-Z. Let i be the unique index such that Z contains at least two vertices of C_i. We have seen that for j>1 V(K_j) is contained in V(C_i) and that V(K_1) intersects V(C_i). Hence every component of G-Z contains the interior of a path of C_i between two vertices of Z. Thus, all three vertices of Z lie on C_i and G-Z has exactly three components, each of which intersects C_i in the interior of a path between two vertices of Z. We note further that each component of G-Z corresponds to a component of G'-Z and hence has an edge to each vertex of Z. Now, if i is not h(k) then K_3 is connected not just in G but also in H because $C_{h(k)}$ is connected and lies in K_3 and all the edges of G-H have both endpoints on $C_{h(K)}$. Furthermore, K_1 and K_2 are components of H-Z. But now contracting these three components into vertices we see that H has a $K_{3,3}$ minor contradicting its minimality. So i is h(k).

Now, we note that the vertices of K_1 and K_2 are joined to none of H-$C_{h(k)}$. Thus, if we let a and b be the vertices of Z which are joined via a path of $C_{h(k)}$ then replacing the path of $C_{h(k)}$ which joins them, uses the third vertex c of Z and is contained in K_1 union K_2 union c by any other such path, would yield a New set of protexting cycles. So no such path can exist. But this is nonsence since all of a,b, and c have a neighbour in both K_1 and K_2, QED.

Question 2: We use the fact discussed in the text that (1) We can read the solution of the instance (G,X) of |X|-realizations of the solution to the instance (G,X') of |X'|-realizations in O(f(|X'|) time provided X' contains X (A partition of X in G if and only if it can be ontained from a realizable partition of X' by deleting the elements of X'-X in each element of the latter partition), and (2) if Z is a subset of X and G-Z splits into l graphs U_1,…,U_l between any two of which there is no edge, then we can read off the solution to the instance (G,X) of |X|-Realizations in O(f(|X|,l)) time from the solution to the realization

instances (G[Z], Z), $(H_1,X_1),\ldots,(H_l,X_l)$ where H_i is the subgraph of G induced by the union of $V(U_i)$ and the vertices of Z with an edge to $V(U_i)$ and X_l is the intersection of $V(H_l)$ and X (From a choice of of a realizable partition of X_i in H_i for each i we can obtain a partition of X which is realizable in G by (i) adding edges so that each element of each of the partitions is a clique, and taking the partition of X given by the vertex sets of the resultant graphs. Conversely every realizable partition of X in G arises in this way. The number of choices we need to consider here is bounded by |X| and k and is independent of |G|).

Applying this second fact when Z is the empty set shows that we can consider each component of G containing a vertex of X separately and then combine the solutions, so we can assume G is connected. We will now use these two facts in a more sophisticated way to show that it is enough to solve k-realizations for all k just on two-connected graphs, and in fact just on three connected graphs without (3,3)-cuts in order to do so in general. .

To begin, we assume we can solve k-realizations for all k on 2-connected graphs and consider some graph G. We consider the block tree of G. If it has only one node then we just apply our algorithm for 2-connected graphs. Otherwise we root it at any cutvertex x. W e traverse the block-tree in post-order, and when we come to a node corresponding to a block B which shares a vertex v with its parent, we delete B-v from G, and this node from the tree if all of its children have been deleted and B-v is disjoint from X. We delete a cutvertex node from the tree if we have deleted all its children. The resultant tree T has at most k leaves. We can assume it is non-empty as otherwise X contains no vertex other than x, in which case the 1-Realization or 0-Realization problem is trivial.

For each cutvertex v of G such that the corresponding node n_v of the block tree was not deleted, we let G_v be the union of the block nodes in the maximal rooted subtree of T rooted at n_v. We let Z_v consist of all of the cutvertices corresponding to the grandchildren of v in T. We note that $|Z_v|$ is at most k. We let X_v be the union of v, and (X intersection $V(G_v)$). We will traverse T in postorder computing for each n_v in turn the solution to the instance (G_v,X_v) of $|X_v|$-realizations. We read the solution to this problem off from the solution to the instance (G_v,X_v union Z_v) of |X_v union Z_v|-realizations.

Apply Fact 1, we see that we need only solve the subproblem described in the statement of the fact for each graph U of G-Z_v such that U either (i) is G[Z], (ii) is G_w-w for some grandchild n_w of n_v and has edges to w but no edges to Z_v-w, or (iii) is B-Z_v for some block which is a child of n_v in T and has edges exactly to those vertices of Z_v in B. In the first case, we are considering a bounded size graph, in the second case we have already solved the subproblem corresponding to U in our postorder traversal. In the third case, we are considering a 2-connected block of G. Since we consider each block at most once, the fact that we can solve k'-realizations on 2-connected graphs in linear time for all $k'<|X|+|Z_v|+1<2k+1$ implies that we can solve it in all graphs in linear time for k.

To prove that it is enough to solve the problem on 3-connected graphs, we apply exactly the same dynamic programming approach to the Lipton-Tarjan decomposition tree of a 2-connected graph. Now however, when in constructing T we delete a block node B because it has become a leaf of the tree and for the cutset (x,y) corresponding to the parent of B we have B-x-y is disjoint from X, we will have to add the edge xy to G. Furthermore, the size of the set of vertices corresponding to a cutnode of the graph and its grandchildren is bounded by 2k+2 rather than k+1.

In exactly the same way, we can use the (3,3)-block tree to show that it is enough to solve k-realizations in linear time for all k in graphs which are planar or L. Again, when deleting a block B corresponds to a node which has become a leaf and such that for the cut Z corresponding to its parent in the tree, B-Z is disjoint from X, we will add edges to make Z a 3-clique.

Since we can solve k-realizations in linear time on planar graph and |L| is 8, the desired result follows.

Question 3: We will set f(l) to be 2g(2(l choose 2)-3, 2(l choose 2)) +6. We will use the edges off the k by k grid as the edges of a K_l model and then apply our results on routing in the plane to the unique embedding of the grid (with the unique non-square face as the infinite face), to find the trees of the model. We can assume l is at least 5 because K_4 is a minor of the 3 by 3 grid.

We consider a minimal counterexample to the problem. This implies that the first row, first column, kth row, and kth column of the grid all contain an endpoint of an off-grid edge. By hypothesis they both contain at most one such endpoint. We let X* be the set of at most four such endpoints lying in these rows and columns.

We arbitrarily choose a bijection f from the edges of K_l to the edges off the grid. For each edge uv of K_l we choose one endpoint of f(uv) to correspond to u, this is h(u,uv). The other endpoint of f(uv) corresponds to v and is h(v,uv). We make this choice arbitrarily, except that if we have that for some u,v,w the vertex on the top row of the grid is h(u,uv) and the vertex on the bottom row is h(u,uw) then we swap the choice we made for this second edge (leaving all other choices the same). We set X to be the union of all the endpoints of all the non-grid edges. We consider the partition Δ of X into l sets, indexed by the vertices of K_l. The element of Δ corresponding to u consists of the l vertices which are h(u,e) for some edge e with u as an endpoint. To prove that G has a K_l model it remains to show that Δ is realizable in the grid.

To this end for each vertex v of X-X*, we cut out a small disc intersecting the drawing of the grid only at v. We also cut out a cuff C* whose interior is contained in the infinite face which intersect G precisely in X*. , So the grid is now embedded in an (2(l choose 2)-3)-punctured plane $\sum$. This will allow us to use Theorem ** to show that Δ is realizable in the grid.

To begin we must show that Δ is realizable in $\sum$. I.e. that there is a planar embedding of the realizability graph which respects the clockwise orientation around each cuff vertex. Since every cuff vertex has degree 1 except for that corresponding to the cuff in the infinite face, it is enough to show that the realizability graph is planar, as we can always choose clockwise so that it agrees with the cyclic order of the edges around one vertex. Now the realizability graph is obtained from a forest F of trees by adding a cycle C of length four through four leaves in at least two trees and then adding a vertex x adjacent to the vertices of C. We suppose that the realizability graph is non-planar and consider the smallest subforest F' of F such that adding C and x to F' yields a non-planar graph F*. It will have been formed in the same way. No tree of F has a leaf off C, as otherwise we could delete it without affecting planarity. So, there are at most four leaves and two trees. If one of the trees has one leaf then it is a vertex, and it is easy to see that F* is planar. If both leaves have two trees then they are paths and the

graph will be planar unless both paths join antipodal points of C. But by our clever choice of H, this is impossible.

Thus, to show that Δ is realizable in the grid, we need simply show that the schisms, O-arcs, and I-arcs are long enough. In doing so, we let C_1 be the unique cycle of G formed by subpath of rows f(l)/2-2 and k-f(l)/2+2 and columns f(l)/2-2 and k-f(l)/2+2 and let C_2 be the unique cycle of G formed by subpath of rows f(l) and k-f(l) and columns f(l) and k-f(l). For each vertex x in X-X*, we let C_x be the column of the grid containing x.

Now for every face of the grid of length four, the indices of the vertices on that face differ by at most 1. Thus for any schism, I-arc, or O-arc J of length at most f(l)/2-3 we have: (I) if J does not intersect the infinite face, both the set of row indices and the set of column indices of the vertices of J form a consecutive sequence, and (II) if J does intersect the infinite face then every vertex on J has an index which is at most f(l)/2-3 or at least k-f(l)/2-3.

It follows that no linking I-arc can have length less than f(l)/2 -3. Furthermore, every O-arc, schism J of length less than f(l)/2-3, or non-planar looping I-arc is either (i) contained in the infinite face of the embedding of C_2 which is a subembedding of our embedding of the grid, or (ii) is in the non-infinite face of the embedding of C_1 which is subembedding of our embedding of the grid.
In the first case, the intersection of $\sum$ with the non-infinite face of C_2 lies in one component of $\sum$-J. Thus this component contains all the cuffs except possibly C*. So J must be an O-arc which is not a schism, or a near-planar I-arc with its endpoints on C*. In the second case, C_1 and C* are contained in one component of $\sum$-J If the cuff containing x is separated from C* by J, then J must intersect C_x. By the consecutiveness of the indices of the vertices of J, we see that this can occur for at most one cuff and that if it does occur for a cuff then J cannot have endpoints on any other cuff. . So J must be an O-arc which is not a schism, or a near planar looping I-arc.

So, it remains to show that if for some cuff C, J is a near-planar looping I-arc with endpoints on C or an O-arc surrounding C then J satisfies the hypothesis of Theorem *.*.

If C is not C* then if J surrounds the cuff contiaing X or cuts off a cuff component containing x then it must intersect C_x and hence contains a vertex and has the required length. If C is C* then J has the required length unless it has length three or less. Since X-X* has at least 16 elements, this implies that there is some x in X-X* such that J is disjoint from C_x. So, if J cuts off a vertex v of X from the cuffs other than C* and thus from x and C_x , then J must contain a vertex in the row of the grid which contains v. By the distinctness of the indices, this vertex is not in X, so if J is a looping I-arc then it is on the interior of J. Thus, by the distinctness of the indices again, J satisfies the hypotheses of Theorem *.* QED,