Stop Minding Your P's and Q's: Simplified Planarity by Edge Addition

John M. Boyer
PureEdge Solutions Inc.
jboyer@PureEdge.com; jboyer@acm.org

Wendy Myrvold *
University of Victoria
wendym@cs.UVic.ca

September 30, 2003

Abstract

A graph is *planar* if it can be drawn on a plane without any crossing edges. This paper presents a new $O(n)$ planar embedding method that avoids some of the complexities of previous approaches. Our new algorithm easily permits the extraction of a planar obstruction (a subgraph homeomorphic to $K_{3,3}$ or K_5) in $O(n)$ time if the graph is not planar. Our method was inspired by observations of the behavior of the PQ-tree method of Booth and Lueker. The P-nodes in a PQ-tree help to represent parts of the partially embedded graph that can be permuted, and the Q-nodes represent parts that can be flipped. We avoid Q nodes by relaxing the notion of vertex orientation so that biconnected components can be flipped in $O(1)$ time. We avoid the use of P-nodes by not connecting subgraphs together until they become biconnected, and we exploit the fact that this can occur when a single edge is added to the partial embedding.

1 Introduction

A *graph* G contains a set V of vertices and a set E of edges, each which corresponds to a pair of vertices from V. Throughout this paper, n denotes the number of vertices of a graph and m indicates the number of edges. A *planar drawing* of a graph is a rendition of the graph on a plane with the vertices at distinct locations and no edge intersections (except at their common vertex endpoints). A graph is *planar* if it admits a planar drawing, and a *planar embedding* is an equivalence class of planar drawings described by the clockwise order of the neighbors of each vertex [9]. In this paper, we focus on a new $O(n)$ planar embedding method. Generating a planar drawing is often viewed as a separate problem, in part because drawing algorithms tend to create a planar embedding

*Supported by NSERC.

as a first step and in part because drawing can be application-dependent. For example, our notion of a suitable rendering for a graph may differ substantially if the graph represents an electronic circuit versus a hypertext book.

We assume the reader has some familiarity with basic graph theory appearing for example in [11, 26], including depth first search (DFS), the adjacency list representation of graphs, and the rationale for focusing on undirected graphs with no loops or parallel edges. We also assume knowledge of basic definitions related to planarity, such as those for proper face, external face, cut vertex and biconnected component. Because a planar graph on more than two vertices has $3n-6$ or fewer edges, the input graph is restricted to at most $3n-5$ edges, enough for all planar graphs and to find a minimal subgraph obstructing planarity in any non-planar graph.

Kuratowski proved that a non-planar graph must contain a subgraph *homeomorphic* to either K_5 or $K_{3,3}$ (subgraphs in the form of K_5 (Figure 1(a)) or $K_{3,3}$ (Figure 1(b)) except that paths can appear in place of the edges). Just as a planar embedding provides a simple certificate to verify a graph's planarity, the indication of a *Kuratowski subgraph* (a subgraph homeomorphic to K_5 or $K_{3,3}$) provides a simple certificate of non-planarity. In some applications, finding a Kuratowski subgraph is a first step in eliminating crossing edges in the graph. For example, in a graph representing an integrated circuit, an edge intersection would be indicative of a short-circuit, which could be repaired by replacing the crossed edge in an identified Kuratowski subgraph with a subcircuit of exclusive-or gates [17].

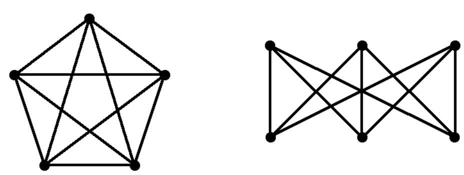

(a) K_5 **(b) $K_{3,3}$**

Figure 1: The Planar Obstructions K_5 and $K_{3,3}$

The first $O(n)$ planarity test algorithm is due to Hopcroft and Tarjan [12]. The method first embeds a cycle C of a biconnected graph, then it breaks the remainder of the graph into a sequence of paths that can be added either to the inside or outside of C. Since a path is added to the partial embedding when it is determined that planarity can be maintained, this method has sometimes been called the *path addition* method. Despite the abundance of additional information on this algorithm [8, 19, 21, 27, 28], the method is known to be complex (e.g. [6, p. 55]), and its implementation in LEDA appears to be slower other linear time algorithms implemented in LEDA [4].

The other well-known method of planarity testing proven to achieve linear time began with an $O(n^2)$ algorithm due to Lempel, Even and Cederbaum [15].

The algorithm begins by creating an s,t-numbering for a biconnected graph. One property of an s,t-numbering is that there is a path of higher numbered vertices leading from every vertex to the final vertex t, which has the highest number. Thus, there must exist an embedding $\tilde{G}_k$ of the first k vertices such that the remaining vertices ($k+1$ to t) can be embedded in a single face of $\tilde{G}_k$. This planarity algorithm was optimized to linear time by a pair of contributions. Even and Tarjan [10] optimized s,t-numbering to linear time, while Booth and Lueker [1] developed the PQ-tree data structure, which allows the planarity test to efficiently maintain information about the portions of the graph that can be permuted or flipped before and after embedding each vertex. Since the permutations and flips of a PQ-tree are performed to determine whether a vertex can be added to the partial embedding while maintaining planarity, this method has become known as a *vertex addition* method.

Chiba, Nishizeki, Abe and Ozawa [6] augmented the PQ-tree operations so that a planar embedding is computed as the PQ-tree operations are performed. Achieving linear time with the vertex addition method is also quite complex [14], partly because many PQ-tree templates are left for the reader to derive [1, p. 362]. There are also non-trivial rules to restrict processing to a 'pertinent' subtree of the PQ-tree, prune the pertinent subtree, and increase the efficiency of selecting and applying templates (since more than one is often applied during the processing for one vertex).

An inspired planarity characterization by de Fraysseix and Rosenstiehl [7] leads to $O(n)$ planarity algorithms, though the paper does not develop the linear time methodology. However, the planarity characterization is particularly interesting because it is the first to examine planarity in terms of conflicts between back edges as seen from a bottom-up view of the depth first search tree.

In 1999, a pair of new vertex addition algorithms were presented that also characterized planarity in terms of conflicts between back edges as seen from a bottom-up view of the depth first search tree. Both also exploited the relationship between cut vertices, biconnected components and depth first search that were originally presented by Tarjan [24]. The expressed purpose of these works was to simplify linear time planarity. The first of these works was the early formulation [3] of the planarity algorithm reported in this paper, a LEDA implementation of which was described in [4]. First, the vertices and depth first search (DFS) tree edges are placed into a partial embedding. Then, each vertex v is processed in reverse order of the depth first indices. The back edges from v to its DFS descendants are added to the partial embedding if it is determined that they can all be added while preserving planarity in the partial embedding.

Later in 1999, Shih and Hsu presented the PC-tree [23], which represented a significant reformulation [25] of their earlier planarity algorithm for triconnected graphs [22] to cover the planarity problem for all biconnected graphs. The PC-tree is similar to a PQ-tree except the Q-nodes are replaced by C-nodes and the operations are simpler. The PC-tree exploits some of the same properties of graphs as our earlier method [3], but due to some errors in [23] first reported in [2], the PC-tree formulation in [23] does not recognize some non-planarity conditions pertaining to C-nodes and does not correctly prove how a particular

C-node condition implies non-planarity. In a book chapter in press [13], Hsu and McConnell correct a number of the C-node problems, but introduce a new one that causes the method to fail on some K_5 homeomorphs. The corrections appearing in [5] describe the considerations that were necessary to produce a correct, linear time PC-tree implementation reported therein (and in [20]).

The method we present in this paper avoids many of the complexities of prior approaches. We present the essential details of our original formulation [3], but we also improve the method by not testing for non-planarity conditions before embedding all the back edges for a vertex. This eliminates the difficulties that have lead to problems for PC-trees as well as the numerosity of non-planarity detection rules used by an implementation [4] of our prior method.

Our new formulation instead adds each back edge to the partial embedding as soon as it is determined that the edge can be added while maintaining planarity, and we therefore refer to the new formulation as an *edge addition* method. Note that edge addition does not mean the edges are added at random. They are added in a particular order to exploit other mathematical properties that ensure correctness and efficiency, which is analogous to the orders imposed on vertices and paths by vertex addition and path addition methods. However, we appear to achieve greater simplicity and efficiency from edge addition since planar embedding is, after all, a matter of avoiding edge crossings if possible.

Section 2 describes the partial embedding data structure as a collection of the biconnected components that develop as edges are added, and it presents the fundamental operation of adding an edge to the partial embedding. Since an edge may biconnect previously separated biconnected components, they are merged together so that a single biconnected component results from adding the new edge. Section 3 explains the key constraint on the fundamental operation of adding an edge, which is that the external face must contain all vertices that will be involved in further edge additions. Due to this constraint, some biconnected components may need to be flipped before they are merged, and Section 4 describes how our method flips a biconnected component in constant time by relaxing the consistency of vertex orientation in the partial embedding.

Based on these principles, Section 5 presents our planar embedder. Section 6 proves that our planarity algorithm achieves linear time performance, and Section 7 proves that it correctly distinguishes between planar and non-planar graphs. In turn, the proof provides a constructive basis for presenting our new Kuratowski subgraph isolator, which appears in Section 8. Finally, Section 9 presents concluding remarks.

2 The Fundamental Operation: Edge Addition

The planarity algorithm described in this paper adds each edge of the input graph G to an embedding data structure $\tilde{G}$ that maintains the set of biconnected components that develop as each edge is added. As each new edge is embedded in $\tilde{G}$, it is possible that two or more biconnected components will be merged together to form a single, larger biconnected component. Figure 2 illustrates the graph theoretic basis for this strategy. In Figure 2(a), we see a connected graph that contains a cut vertex v whose removal, along with its incident edges, separates the graph into the two connected components shown in Figure 2(b). Thus, the graph in Figure 2(a) is represented in $\tilde{G}$ as the two biconnected components shown in Figure 2(c). Observe that the cut vertex v is represented in each biconnected component that contains it. Observe that the addition of a single edge (u, w) with endpoints in the two biconnected components results in the single biconnected component depicted in Figure 2(d). Since v is no longer a cut vertex, the representation no longer requires an extra vertex to represent v.

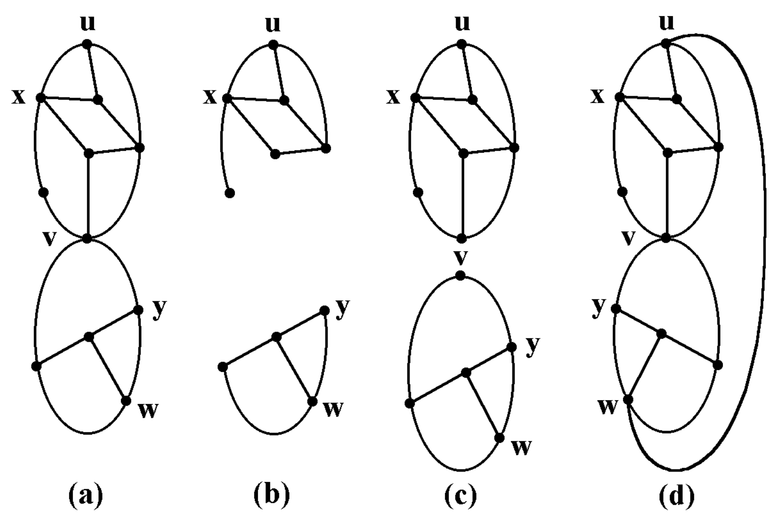

Figure 2: (a) A Cut Vertex v, (b) Removing v results in more connected components, (c) The biconnected components separable by v, (d) When edge (u, w) is added, v is no longer a cut vertex (by flipping the lower biconnected component, y remains on the external face)

Indeed, Figure 2(d) illustrates the fundamental operation of the edge addition planarity algorithm. A single edge biconnects previously separable biconnected components, so these are merged together when the edge is embedded, resulting in a single larger biconnected component B. Moreover, the key constraint on this edge addition operation is that any vertex in B must remain on the external face of B if it must be involved in the future embedding of an edge. Hence, a biconnected component may need to be flipped before it is merged. For example, the lower biconnected component in Figure 2(d) was merged, but it was also flipped on the vertical axis from v to w so that y has remained on the external face.

3 External Activity

One of the key problems to be solved in an efficient planarity algorithm is how to add some portion of the input graph to the embedding in such a way that little or no further adjustment to the embedding is necessary to continue adding more of the input graph to the embedding. As described in Section 1, the PQ-tree method exploits a property of an s,t-numbered graph that allows it to create an embedding $\tilde{G}_k$ of the first k vertices such that the remaining vertices ($k+1$ to t) can be embedded in a single face of $\tilde{G}_k$. We observe that an analogous property exists for a depth first search tree: Each vertex has a path of lower numbered DFS ancestors that lead to the DFS tree root. Therefore, our method processes the vertices in reverse order of their depth first indices (DFI) so that every unprocessed vertex has a path of unprocessed DFS ancestors leading to the last vertex to be processed, the DFS tree root. As a result, all unprocessed vertices must appear in a single face of the partial embedding $\tilde{G}$, and the flipping operations described in the prior section are designed to allow that face to be the common external face shared by all biconnected components in $\tilde{G}$.

During the fundamental operation of adding an edge, we must be able to determine when to flip a biconnected component so that vertices with unembedded edge connections to unprocessed vertices remain on the external face. For each vertex v in reverse DFI order, our method first embeds the tree edges from v to its children, then the back edges from v to its descendants. For any descendant w of v, we define a vertex w in a biconnected component B to be *externally active* if there is a path from w to a DFS ancestor u of v consisting of a back edge plus zero or more DFS descendants of w, none of which are in B. Thus, an externally active vertex w will be involved in the future embedding of edges after the processing of v, either as the descendant endpoint of a back edge or as a cut vertex that will be a biconnected component merge point during the embedding of a back edge to some descendant of w.

The external activity of each vertex can be efficiently maintained as follows. The *lowpoint* of a vertex is the DFS ancestor of least DFI that can be reached by a path of zero or more descendant DFS tree edges plus one back edge, and it can be computed in linear time by a post-order traversal of the depth first search tree [24]. During preprocessing, we first obtain the *least ancestor* directly adjacent to each vertex by a back edge, then we compute the *lowpoint* of each vertex. Next, we equip each vertex with a list called *separatedDFSChildList*, which initially contains references to all DFS children of the vertex, sorted by their lowpoint values. To do this in linear time, categorize the vertices into ascending order by their lowpoint values, then add each to the end of the separatedDFSChildList of its DFS parent. To facilitate constant time deletion from a separatedDFSChildList, it is made circular and doubly linked, and each vertex is also equipped with a parameter that indicates the representative node for the vertex in the separatedDFSChildList of its DFS parent. When a biconnected component containing a cut vertex w and one if its children c is merged with the biconnected component containing w and the DFS parent of w, then the representative of c is deleted from the separatedDFSChildList of w. Thus, a vertex w is externally active during the processing of v if it either has a least ancestor less than v or the first element in its separatedDFSChildList has a lowpoint less than v.

4 Flipping in Constant Time

Section 2 described the fundamental operation of adding a back edge to $\tilde{G}$ that might biconnect previously separated biconnected components in $\tilde{G}$. This necessitated merging the biconnected components together as part of adding the new edge to $\tilde{G}$. Then, Section 3 discussed the key constraint on the biconnected component merging process, which was that externally active vertices had to remain on the external face of the new biconnected component formed from the new edge and the previously separated biconnected components. This necessitated flipping some of the biconnected components.

The easiest method for flipping a biconnected component is to simply invert the adjacency list order, or *orientation*, of each of its vertices. However, it is easy to create graphs in which $\Omega(n)$ vertices would be inverted $\Omega(n)$ times during the embedding process. To solve this problem, we first observe that a biconnected component B in which vertex r has the least depth first index (DFI) never contains more than one DFS child of r (otherwise, B would not be biconnected since depth first search could find no path between the children except through r). We say that r is the *root* of B, and the DFS tree edge connecting the root to its only DFS child c in B is called the *root edge* of B. We denote a root vertex r^c to distinguish it from all other copies of r in $\tilde{G}$. Next, we observe that there are at most $O(n)$ biconnected component merge operations since each biconnected component root is only merged once and is associated with one DFS child of its corresponding non-root vertex. Thirdly, we observe that a flip operation can only occur immediately before a merge. Therefore, a strategy that achieves constant time performance per flip operation will cost linear time in total. Finally, we observe that it is only a little more difficult to traverse the external face of a biconnected component if some vertices have a clockwise orientation and others have a counterclockwise orientation. Therefore, we flip a biconnected component by inverting the orientation of its root vertex.

A planar embedding with a consistent vertex orientation can be recovered in post-processing if an additional piece of information is maintained during embedding. We equip each edge with a *sign* initialized to $+1$. When a biconnected component must be flipped, we only invert the adjacency list orientation of the root vertex r^c so that it matches the orientation of the vertex r with which it will be merged. Then, the sign of the root edge (r^c, c) is changed to -1 to signify that all vertices in the DFS subtree rooted by c now have an inverse orientation. Since all of the edges processed by this operation were incident to a root vertex beforehand and a non-root vertex afterward, only constant work per edge is performed by all merge and flip operations during the embedding process. Moreover, a planar embedding for any biconnected component can be recovered at any time by imposing the orientation of the biconnected component root on all vertices in the biconnected component. If the product of the signs along the tree path from a vertex to the biconnected component root is -1, then the adjacency list of the vertex should be inverted. This can be done with a cost commensurate with the size of the biconnected component by using a depth first search over the existing tree edges in the biconnected component.

Figure 3 helps to illustrate the biconnected component flip operation. Each vertex has a black dot and a white dot that signify the two pointers to the edges that attach the vertex to the external face (if it is on the external face). Observe the dots of each vertex to see changes of vertex orientation. Vertices 2 and 6 are square to signify that they are externally active due to unembedded back edges to vertex 0. The goal is to embed the edge $(1,4)$.

In Figure 3(a), consistently following the black dot pointers yields a counterclockwise traversal of the external face of any biconnected component. In Figure 3(b), the biconnected component with root 3^4 is flipped so that edge $(1,4)$ can be embedded along the left side while the externally active vertices 2 and 6 remain on the external face. Note that the orientations of vertices 4, 5 and 6 did not actually change relative to Figure 3(a). For example, the black dot in vertex 5 still leads to vertex 4, and the white dot in vertex 5 still leads to vertex 6. The result of embedding edge $(1,4)$ appears in Figure 3(c).

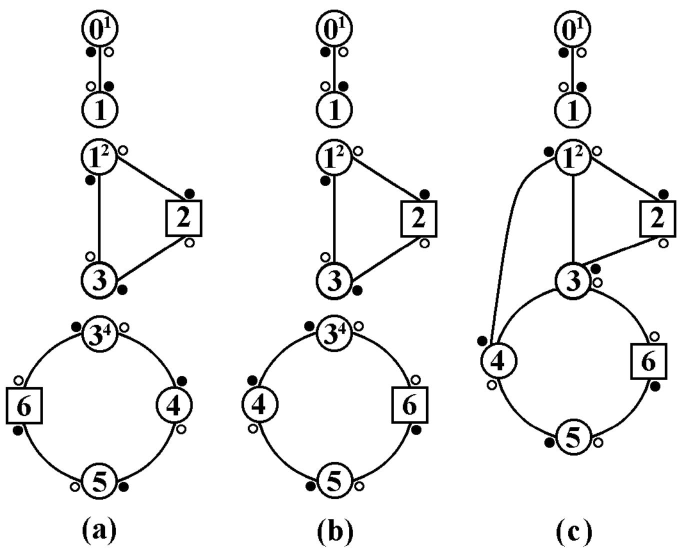

Figure 3: (a) $\tilde{G}$ before embedding edge $(1,4)$. Vertices 2 and 6 are square to indicate they are externally active. (b) The biconnected component rooted by 3^4 is flipped, without altering the orientations of vertices 4 to 6, whose black and white dots indicate the same edges as they did before the flip. (c) Edge $(1,4)$ is embedded, and vertices 2 and 6 have remained on the external face. The orientations of vertices 4 to 6 are not consistent with those of vertices 1 to 3, but the external face can still be traversed by exiting each vertex using whichever edge was not used to enter it.

Finally, we observe that the only new problem introduced by this technique is that extra care must be taken to properly traverse the external face of a biconnected component. In Figure 3(c), it is clear that following the black dots consistently no longer corresponds to a counterclockwise traversal of the external face of the biconnected component rooted by 1^2. The black and white dots signify links from a vertex to nodes of its adjacency list. A typical adjacency list representation has only one link or pointer from a vertex to a node its its adjacency list. In $\tilde{G}$, we use two link pointers so that we may indicate the nodes representing both edges that hold a vertex on the external face (if indeed the vertex is on the external face). These links can be traversed to obtain the edge leading to the next neighbor along the external face. However, the order of the links to the external face edges is reflective of the orientation of a vertex, and the orientation of vertices in a biconnected component can vary between clockwise and counterclockwise. Hence, whenever our method traverses an external face, it keeps track of the vertex w being visited but also the link to the edge that was used to enter w, and the opposing link is used to obtain the edge leading to the successor of w on the external face.

5 Planar Embedding by Edge Addition

Based on the key operations and definitions of the prior sections, this section presents our planar embedding algorithm. Given an input graph G, we begin with preprocessing steps to identify a depth-first search tree and the corresponding back edges, to compute the least ancestor and lowpoint values, and to initialize the embedding structure $\tilde{G}$ so that it can maintain the notion of external activity as defined in Section 3. Then, each vertex v is processed in reverse DFI order to add the edges from v to its descendants.

The DFS tree edges from v to its children are added first to $\tilde{G}$. As a result, for each child c of v, $\tilde{G}$ will contain a new biconnected component consisting solely of the edge (v^c, c). Although it is not necessary in an implementation, each biconnected component containing only a tree edge can be regarded as if it contained two parallel edges between the root v^c and the child c so that the external face forms a cycle, as is the case for all larger biconnected components in $\tilde{G}$.

Next, the back edges from v to its descendants must be added. If we were to systematically perform a depth first search of the DFS subtree rooted by v to find all descendant endpoints of back edges to v, then clearly there are graphs for which $\Omega(n^2)$ performance would result. Therefore, we must restrict the processing to only the *pertinent subgraph* which is the set of biconnected components in $\tilde{G}$ that will be merged together due to the addition of new back edges to v. Our method identifies the pertinent subgraph with the aid of a routine we call the Walkup, which is discussed in Section 5.1. Once the pertinent subgraph is identified, our method adds the back edges from v to its descendants in the pertinent subgraph, while merging and flipping biconnected components as necessary to maintain planarity in $\tilde{G}$ and keeping all externally active vertices on the external faces of the biconnected components embedded in $\tilde{G}$. This phase is performed with the aid of a routine we call the Walkdown, which is discussed in Section 5.2. Figure 4 presents pseudocode for our planarity algorithm.

As will be shown in the proof of correctness in Section 7, the Walkdown only fails to embed all back edges from v to its descendants if the input graph is non-

Figure 4: The Edge Addition Planarity Algorithm

Procedure: Planarity
input: Simple undirected graph G with $n \geq 2$ vertices and $m \leq 3n - 5$ edges
output: PLANAR and an embedding in $\tilde{G}$, or
NONPLANAR and a Kuratowski subgraph of G in $\tilde{G}$

(1) Perform depth first search and lowpoint calculations for G
(2) Create and initialize $\tilde{G}$ based on G, including creation of separatedDFSChildList for each vertex, sorted by child lowpoint

(3) For each vertex v from $n - 1$ down to 0

(4) for each DFS child c of v in G
(5) Embed tree edge (v^c, c) as a biconnected component in $\tilde{G}$

(6) for each back edge of G incident to v and a descendant w
(7) Walkup($\tilde{G}$, v, w)

(8) for each DFS child c of v in G
(9) Walkdown($\tilde{G}$, v^c)

(10) for each back edge of G incident to v and a descendant w
(11) if $(v^c, w) \notin \tilde{G}$
(12) IsolateKuratowskiSubgraph($\tilde{G}$, G, v)
(13) return (NONPLANAR, $\tilde{G}$)

(14) RecoverPlanarEmbedding($\tilde{G}$)
(15) return (PLANAR, $\tilde{G}$)

planar, so a routine to isolate a Kuratowski subgraph is invoked. A discussion of how to implement this routine appears in Section 8. Otherwise, if all tree edges and back edges are embedded, then the planar embedding is recovered as discussed in Section 4.

Most of the remainder of this section discusses the Walkup and Walkdown. However, in order to understand their processing, we must make a few more definitions, mostly to provide a context for the notion of a pertinent subgraph in terms of our embedding structure $\tilde{G}$, which manages a collection of biconnected components. A biconnected component with root w^c is a *child biconnected component* of w. A non-root vertex w is *pertinent* if G has a back edge (v, w) not embedded in $\tilde{G}$ or w has a child biconnected component in $\tilde{G}$ that contains a pertinent vertex. A *pertinent biconnected component* contains a pertinent vertex. An *externally active biconnected component* contains a non-root externally active vertex. A non-root vertex is *internally active* if it is pertinent but not externally active, and an *internally active biconnected component* is pertinent but not externally active. A vertex is *inactive* if it is neither internally nor externally active.

5.1 The Walkup

In this section, we discuss a subroutine called Walkup, which is invoked by the core planarity algorithm once for each back edge (v, w) to help gather information needed by the Walkdown to embed back edges from v to its descendants. Specifically, the purpose of the Walkup is to identify vertices and biconnected components that are pertinent due to the given back edge (v, w). Vertex w is pertinent because it is the direct descendant endpoint of a back edge to be embedded. Each cut vertex in $\tilde{G}$ along the DFS tree path strictly between v and w, denoted $T_{v,w}$, is pertinent because the cut vertex will become a merge point during the embedding of a back edge.

To help us mark as pertinent the descendant endpoints of back edges to v, we equip each vertex with a flag called *backedgeFlag*, which is initially cleared during preprocessing. The first action of the Walkup for the back edge (v, w) is to raise the backedgeFlag flag of w. Later, when the Walkdown embeds the back edge (v, w), the backedgeFlag flag is cleared. To help us mark as pertinent the cut vertices in $\tilde{G}$ along the DFS tree path $T_{v,w}$ between v and w, we equip the vertices with an initially empty list called *pertinentRoots*. Let r be a cut vertex in $\tilde{G}$ along $T_{v,w}$ with a DFS child s also in $T_{v,w}$. Then, r^s is a biconnected component root along $T_{v,w}$ in $\tilde{G}$. When the Walkup finds r^s, it adds r^s to the pertinentRoots of r. Later, when the Walkdown merges r and r^s, then r^s is removed from the pertinentRoots list of r. Thus, a vertex is determined to be pertinent if its backedgeFlag flag is set or if its pertinentRoots list is non-empty.

Although not strictly necessary, the Walkdown implementation benefits if the Walkup places the roots of externally active biconnected components after the roots of internally active biconnected components. Thus, a biconnected component root r^s is prepended to the pertinentRoots list of r unless r^s is the root of an externally active biconnected component (i.e. if $\text{lowpoint}(s) < v$), in which case it is appended instead.

A few strategies must be employed to efficiently implement the traversal phase of the Walkup that sets the pertinentRoots lists of cut vertices in $\tilde{G}$ along $T_{v,w}$. The goal is to ensure that the work done by all invocations of Walkup in step v is commensurate with the sum of the sizes of the proper faces that will be formed when the back edges from v to its descendants are added to $\tilde{G}$. Therefore, we do not simply traverse the tree path $T_{v,w}$ in $\tilde{G}$ looking for biconnected component roots. Instead, our method traverses the external face of each biconnected component that contains part of $T_{v,w}$ to find the root. If w is the point of entry in a biconnected component, there are two possible external face paths that can be traversed to obtain the root r^s. Since the path must become part of a proper face, it cannot contain an externally active vertex between w and r^s. However, it is not known which path is longer, nor if either contains an externally active vertex. Therefore, both paths are traversed in parallel to ensure that r^s is found with a cost no greater than twice the length of the shorter external face path. Once r^s is located, it is added to the pertinentRoots of r, then r is treated as the entry point of the biconnected component, and the Walkup reiterates until $r^s = v^c$.

By traversing external face paths in parallel, rather than simply searching the tree path $T_{v,w}$, a Walkup can find the roots of biconnected components made pertinent by the back edge (v, w). However, the intersection of T_{v,w_1} and T_{v,w_2} for two back edges (v, w_1) and (v, w_2) can be arbitrarily large, so the Walkup must be able to detect when it has already identified the roots of pertinent biconnected components along a path due to a prior invocation for another back edge of v. To solve this problem, we equip each vertex with a flag called *visited*, which is initially cleared. Before visiting any vertex, the Walkup checks whether the visited flag is set, and ends the invocation if so. Otherwise, after visiting the vertex, the Walkup sets the visited flag and adds the vertex to a list of vertices whose visited flags must be cleared at the end of step v.

5.2 The Walkdown

In this section, we discuss a subroutine called Walkdown, which is invoked by the core planarity algorithm once for each DFS child c of the vertex v to embed the back edges from v to descendants of c. The Walkdown receives the root v^c of a biconnected component B. The Walkdown descends from v^c, traversing along the external face paths of biconnected components in the pertinent subgraph (identified by the Walkup described in Section 5.1). When the Walkdown encounters the descendant endpoint w of a back edge to v, the biconnected components visited since the last back edge embedding are merged into B, and the edge (v^c, w) is then added to B. As described in Section 4, each biconnected components that is merged into B may be flipped if necessary to ensure that externally active vertices remain on the external face of B. The external face paths traversed by the Walkdown become part of the new proper face formed by adding (v^c, w).

To embed the back edges, the Walkdown performs two traversals of the external face of B, corresponding to the two opposing external face paths emanating from v^c. The traversals perform the same operations and are terminated by the same types of conditions, so the method of traversal will only be described once.

A traversal begins at v^c and proceeds in a given direction from vertex to vertex along the external face in search of the descendant endpoint of a back edge. Whenever a vertex is found to have a pertinent child biconnected component, the Walkdown descends to its root and proceeds with the search. Once the descendant endpoint w of a back edge is found, the biconnected component roots visited along the way must be merged (and the biconnected components flipped as necessary) before the back edge (v^c, w) is embedded. A Walkdown traversal uses an initially empty *merge stack* to help keep track of the biconnected component roots to which it has descended as well as information that helps determine whether each biconnected component must be flipped when it is merged.

A biconnected component must be flipped if the direction of traversal upon entering a cut vertex r changes when the traversal exits the root r^s. For example, in Figure 3(a), v^c is 1^2, and the first traversal exits 1^2 on the black dot link and enters vertex 3 on the white dot link. The traversal then descends to the root 3^4, where it must exit using the white dot link to avoid the externally active vertex 6. Since entry and exit are inverse operations, opposite colored links must be used to avoid a flip. When the same link color is used for both, the direction of traversal has changed and a flip is required. Note that once the orientation of the root vertex 3^4 is inverted in Figure 3(b), the traversal exits from the black dot link to reach vertex 4. As the above example shows, in order to be able to merge and possibly flip a biconnected component, the following pieces of information must have been stored in the merge stack when a Walkdown traversal descended from a vertex r to the root of one of its pertinent child biconnected components: the identity of the root r^s (from which the non-root r is easily calculated), the direction of entry into r, and the direction of exit from r^s.

A Walkdown traversal terminates either when it returns to v^c or when it encounters a non-pertinent, externally active vertex, which we call a *stopping vertex*. If it were to proceed to embed an edge after passing a stopping vertex, then the stopping vertex would not remain on the external face, but it must because it is externally active. Note that a pertinent externally active vertex may become a stopping vertex once the Walkdown embeds the edges that made the vertex pertinent.

Prior to encountering a stopping vertex, if a vertex w is encountered that has more than one pertinent child biconnected component, then the Walkdown must descend to an internally active child biconnected component if one is available. Note that the Walkdown traverses the entire external face of an internally active child biconnected component and returns to w, but once the Walkdown descends to a pertinent externally active child biconnected component, it will encounter a stopping vertex before returning to w (traversal cannot proceed beyond the first externally active vertex in the child biconnected component). So, to ensure that all back edges that can be embedded while preserving planarity are in fact embedded, the Walkdown first embeds a back edge to w (if needed) and descends to its internally active child biconnected components (if any) before processing its pertinent externally active child biconnected components. If the Walkup uses the strategy from Section 5.1 of always putting externally active child biconnected component roots at the end of the pertinentRoots lists, then the Walkdown can simply pick the first element of a pertinentRoots list.

A similar argument to the one above also governs how the Walkdown chooses a direction from which to exit a biconnected component root r^s to which it has descended. Both external face paths emanating from r^s are searched to find the first internally or externally active vertices x and y in each direction (i.e. inactive vertices are skipped). If both paths lead to internally active vertices, then either path can be selected, but if only one path leads to an internally active vertex, then that path must be selected. If the path to the externally active vertex were selected first, then the traversal would encounter a stopping vertex before reaching the internally active vertex. Similarly, if both paths lead to externally active vertices, then the path to a pertinent vertex must be selected (either path can be selected if both are pertinent). Finally, if both paths lead to externally active vertices and neither is pertinent, then both are stopping vertices. In this case, the entire Walkdown (not just the current traversal) can be immediately terminated due to a non-planarity condition described in Section 7.

With two exceptions, the operations described above allow the total cost of all Walkdown operations to be linear. Traversal from vertex to vertex and queries of external activity and pertinence are constant time operations. Merge and flip operations on biconnected components cost constant time per edge. The data structure updates that maintain external activity and pertinence are constant time additions to the biconnected component merge operation. Finally, adding an edge takes constant time. Thus, the cost of all Walkdown operations performed to add an edge can be associated with the size of the proper face that forms when the new edge is added, except that when Walkdown descends to a biconnected component root, it traverses both external face paths emanating from the root, but only one of those paths becomes part of the proper face formed when the new edge is added. It is possible to construct graphs in which a path of length $\Omega(n)$ is traversed but not selected $\Omega(n)$ times. A second problem is that the cost of traversing the path between the last embedded back edge endpoint and the stopping vertex is not associated with a proper face.

Both of these costs can be bounded to a linear total by the introduction into $\tilde{G}$ of *short-circuit edges*, which are specially marked edges added between v^c and the stopping vertex of each Walkdown traversal (except of course when the traversal is terminated by the above non-planarity condition). The short-circuit edge forms a new proper face, which removes the path from the last back edge endpoint to the stopping vertex from the external face. Moreover, this method reduces to $O(1)$ the cost of choosing the path to proceed along after desdending to a biconnected component root. In step v, the immediate external face neighbors of a pertinent root r^s are active because the interceding inactive vertices are removed from the external face by short-circuit edges embedded in step r. Since only two short-circuit edges are added per biconnected component, and each can be associated with a unique vertex (the DFS child in its root edge), at most $2n$ short-circuit edges are added. Short-circuit edges are also specially marked, so they can be easily removed during the post-processing steps that recover a planar embedding or Kuratowski subgraph of G from $\tilde{G}$.

6 Linear Time Performance

In this section we consider the measures that have been taken to make our planarity algorithm achieve $O(n)$ performance. The result is stated as Theorem 6.1.

Theorem 6.1 *Given a graph G with n vertices, algorithm* Planarity *determines whether G is planar in $O(n)$ time.*

Proof. Preprocessing steps for depth first search, least ancestor and lowpoint calculations are implemented with well-known linear time algorithms. Section 3 described a method that allows linear time initialization of data structures for maintaining external activity, which consists of creating a separatedDFSChildList for each vertex, sorted by the lowpoint values of the children in each list. During the run of the main edge embedding loop, the cost of embedding each tree edge is $O(1)$, resulting in linear time total cost. Section 5.1 describes a method for implementing the Walkup to identify the pertinent subgraph with a cost commensurate with the sizes of the proper faces that will be formed by back edges when they are embedded. Similarly, Section 5.2 describes a method for implementing the Walkdown to embed the back edges within the pertinent subgraph with a cost commensurate with the sizes of the proper faces that will be formed by back edges and short-circuit edges. Since the sum of the degrees of all proper faces is twice the number of edges, a total linear cost is associated with all Walkup and Walkdown operations. The cost of the loops that invoke Walkup and Walkdown are associated with the back edges and tree edges, respectively, for which the functions are invoked, yielding constant cost per edge. At the end of the main loop, the test to determine whether any back edge was not embedded by the Walkdown results in an additional constant cost per back edge, for a total linear cost. □

Corollary 6.2 *Given a planar graph G with n vertices, algorithm* Planarity *produces a planar embedding of G in $O(n)$ time.*

Proof. First, all of the edges of G are added to $\tilde{G}$ by the main loop of Planarity, which takes $O(n)$ time by Theorem 6.1. Then, the embedding is recovered from $\tilde{G}$. The short-circuit edges added to optimize the Walkdown (see Section 5.2) are removed, which takes $O(n)$ time. Then, the vertices in each biconnected component are given a consistent vertex orientation as described in Section 4. This operation can do no worse than invert the adjacency list of every non-root vertex at a cost of constant time per edge, or $O(n)$ in total. Finally, $\tilde{G}$ is searched for any remaining biconnected component roots in $O(n)$ time (there will be more than one if G is not biconnected), and they are merged with their non-root counterparts (without flipping) for a cost commensurate with the sum of the degrees of the roots, which is $O(n)$. □

7 Proof of Correctness

In this section, we prove that the algorithm Planarity described in Section 5 correctly distinguishes between planar and non-planar graphs. It is clear that the algorithm maintains planarity of the biconnected components in $\tilde{G}$ as an invariant during the addition of each edge (see Corollary 7.2). Thus, a graph G is planar if all of its edges are added to $\tilde{G}$, and we focus on showing that if the algorithm fails to embed an edge, then the graph must be non-planar.

For each vertex v, the algorithm first adds to $\tilde{G}$ the tree edges between v and its children without the possibility of failure. Later, the back edges from v to its descendants are added by the routine called Walkdown described in Section 5.2. The main algorithm Planarity invokes the Walkdown once for each DFS child c of v, passing it the root v^c of a biconnected component B in which it must embed the back edges between v^c and descendants of c.

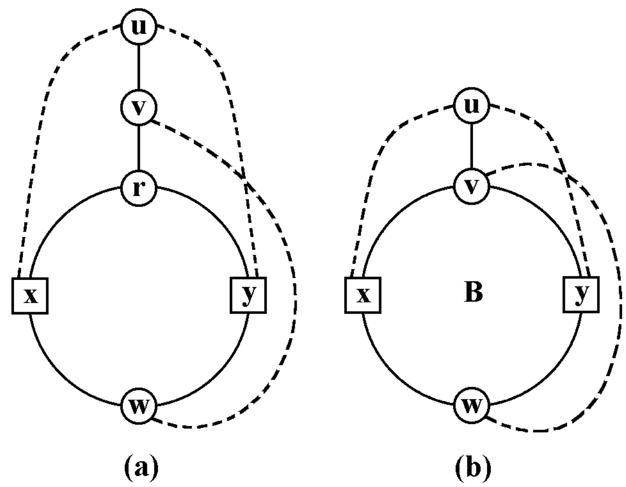

Figure 5: Walkdown Halting Configurations. Squares indicate externally active vertices, and all edges may be paths. Dashed edges include an unembedded back edge. All ancestors of the current vertex v are contracted into u. (a) The Walkdown has descended from the current vertex v to the root r of a biconnected component, but the pertinent vertex cannot be reached along either external face path due to stopping vertices x and y. (b) Each Walkdown traversal from v has encountered a stopping vertex in the biconnected component B that contains a root copy of v. The Walkdown could not reach the pertinent vertex w due to the stopping vertices. As given, this configuration is planar, but Theorem 7.1 proves the existence of additional edges that form one of four non-planar configurations.

For a given biconnected component B rooted by v^c, if the two Walkdown traversals embed all back edges between v and descendants of c, then it is easy to see that B remains planar and the algorithm continues. However, if some of the back edges to descendants of c are not embedded, then we show that the input graph is non-planar. The Walkdown may halt if it encounters two stopping vertices while trying to determine the direction of traversal from a pertinent child biconnected component root, a condition depicted in Figure 5(a). Otherwise, if the Walkdown halts on B without embedding all back edges from v^c to descendants of c, then both Walkdown traversals were terminated by stopping vertices appearing along the external face of B. This condition is depicted by Figure 5(b).

In Figure 5(a), u represents the contraction of all unprocessed ancestors of v, and (u, v) represents the DFS tree path from v to its ancestors. The edge (v, r) represents the path of descent from v to a pertinent child biconnected component rooted by a root copy of r. Square vertices are externally active. The Walkdown traversal is prevented from visiting a pertinent vertex w by stopping vertices x and y on both external face paths emanating from r. The cycle (r, x, w, y, r)

represents the external face of the biconnected component. The dotted edges (u,x), (u,y) and (v,w) represent connections from a descendant (x, y or w) to an ancestor (u or v) consisting of either a single unembedded back edge or a path containing a tree edge to a separated DFS child of the descendant, zero or more additional tree edges, and an unembedded back edge to the ancestor. Similarly, Figure 5(b) shows stopping vertices x and y that prevent traversal from reaching a pertinent vertex w in a biconnected component rooted by a root copy of v.

Both diagrams depict minors of the input graph. Since Figure 5(a) depicts a $K_{3,3}$, the input graph is non-planar. However, Figure 5(b) appears to be planar, so it is natural to ask why the edge (v,w) was not embedded inside B, which the Walkdown could do by embedding (v,w) along the external face, then embedding (v,x) such that (v,w) is inside the bounding cycle of B. In short, there is either some aspect of the connection represented by edge (v,w) or some aspect of the vertices embedded within B that prevents the Walkdown from embedding the connection from w to v inside B. An examination of the possibilities related to these aspects yields four additional non-planarity minors, or five in total, which are depicted in Figure 6. Theorem 7.1 argues the correctness of our algorithm by showing that one of the non-planarity minors must exist if the Walkdown fails to embed a back edge, and the absence of the conditions that give rise to the non-planarity minors contradicts the assumption that the Walkdown failed to embed a back edge.

Theorem 7.1 *Given a biconnected connected component B with root v^c, if the* Walkdown *fails to embed a back edge from v to a descendant of c, then the input graph G is not planar.*

Proof. By contradiction, suppose the input graph is planar but the Walkdown halts without embedding a back edge. To do so, the Walkdown must encounter a stopping vertex. If this occurs because stopping vertices were encountered along both external face paths emanating from the root of a pertinent child biconnected component, then the Walkdown terminates immediately, and the $K_{3,3}$ depicted in Figure 6(a) shows that the input graph is non-planar. Hence, the Walkdown must halt on stopping vertices on the external face of the biconnected component containing v^c.

Figure 6(b) results if the pertinent vertex w has an externally active pertinent child biconnected component. Embedding the connection from a separated descendant of w to v inside B would place an externally active vertex z inside B. Thus, the input graph is non-planar since Figure 6(b) contains a $K_{3,3}$.

Otherwise we consider conditions related to having an obstructing path inside B that contains only internal vertices of B except for two points of attachment along the external face: one along the path $v, \ldots, x, \ldots, w$, and the other along the path $v, \ldots, y, \ldots, w$. The obstructing path, which is called an x-y path, contains neither v nor w. If such an x-y path exists, then the connection from w to v would cross it if the connection were embedded inside B. We use p_x and p_y to denote the points of attachment of the obstructing x-y path.

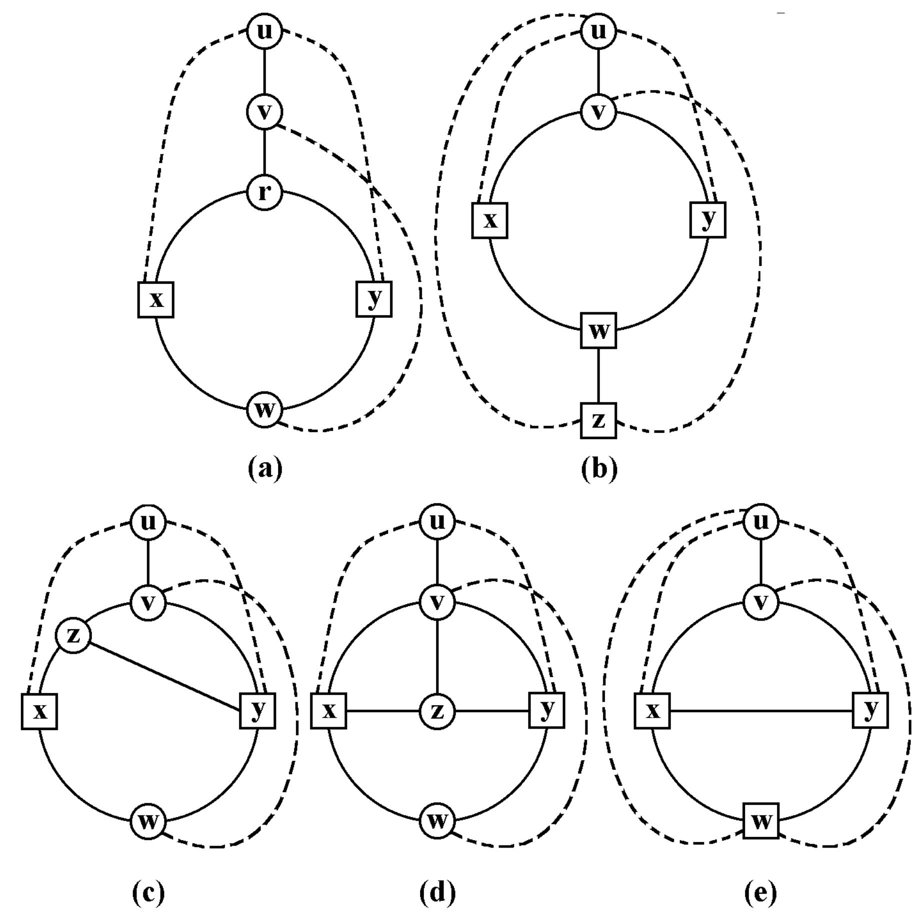

Figure 6: Non-planarity Minors of the Input Graph

In Figure 6(c), the x-y path has p_x attached closer to v than x. Note that p_y can also be attached closer to v than y. In fact, Figure 6(c) also represents the symmetric condition in which p_y is attached closer to v than y (but p_x is attached at x or farther from v than x). In all of these cases, the input graph is non-planar since Figure 6(c) contains a $K_{3,3}$.

In Figure 6(d), a second path of vertices attached to v that (other than v) contains vertices internal to B that lead to an attachment point z along the x-y path. If this second path exists, then input graph is non-planar since Figure 6(d) contains a $K_{3,3}$.

Figure 6(e) depicts the condition of having an externally active vertex (possibly distinct from w) along the lower external face path strictly between p_x and p_y. If this condition occurs, then input graph is non-planar since Figure 6(e) represents a K_5 minor.

Finally, suppose for the purpose of contradiction that the Walkdown failed to embed a back edge and none of the non-planarity conditions described above exist. As mentioned above, due to the absence of the condition of Figure 6(a), the two Walkdown traversals must have ended on stopping vertices along external face paths in the biconnected component B rooted by v^c. By the contradictive assumption, B has a pertinent vertex w along the lower external face path strictly between stopping vertices x and y. We address two cases based on whether or not there is an obstructing x-y path.

If there is no obstructing x-y path, then at the beginning of step v all paths between x and y in $\tilde{G}$ contain w. Thus, w is a DFS ancestor of x or y (or both), and it becomes a merge point when its descendants (x or y or both) are incorporated into B. When the Walkdown first visits w, it embeds a direct back edge from w to v if one is required, so the pertinence of w must be the result of a pertinent child biconnected component. However, the Walkdown preferentially selects and processes internally active child biconnected components of w prior to attaching an externally active pertinent child biconnected component leading to x or y. Thus, the pertinence of w must be due to an externally active child biconnected component, which contradicts the pertinence of w since the condition of Figure 6(b) does not exist.

On the other hand, suppose there is an obstructing x-y path, but non-planarity minors C to E do not apply. The *highest x-y path* is the x-y path that would be contained by a proper face cycle if the internal edges to v^c were removed, along with any resulting separable components. The highest x-y path and the lower external face path from p_x to p_y formed the external face of a biconnected component at the beginning of step v. Let r_1 denote whichever of p_x or p_y was the root of that biconnected component, and let r_2 denote one of p_x or p_y such that $r_1 \neq r_2$. Since the condition of Figure 6(c) does not exist, r_2 is equal to or an ancestor of x or y and was therefore externally active when the Walkdown descended to r_1^s (a root copy of r_1, where s is equal to or a DFS ancestor of r_2). Moreover, the first active vertex along the path that is now the highest x-y path was r_2 because the condition of Figure 6(d) does not exist. Descending from r_1^s along the path that is now the lower external face path between p_x and p_y, the existence of a pertinent vertex w implies that there are no externally active vertices along the path due to the absence of the condition of Figure 6(e). Thus, we reach a contradiction to the pertinence of w since the Walkdown preferentially selects the path of traversal leading from the root of a child biconnected component to an internally active vertex. $\square$

Corollary 7.2 *Given a graph G, algorithm* Planarity *determines whether or not G is planar.*

Proof. For each vertex v in reverse DFI order, the edges from v to its descendants are embedded. The embedding of tree edges cannot fail, and if the Walkdown fails to embed a back edge from v to a descendant, then Theorem 7.1 shows that the graph is not planar. Hence, consider the case in which Planarity embeds all edges from v to its descendants in each step. The edges from v to its ancestors are therefore embedded as those ancestors are processed. When a tree edge is added to $\tilde{G}$, planarity is maintained since the tree edge is added into a biconnected component by itself. When a back edge is added, the preceding merge and flip of biconnected components maintain planarity, and the new back edge is added in the external face region incident to two vertices currently on the external face. Thus, planarity is maintained in $\tilde{G}$ for all edges added, so G is planar if all of its edges are added to $\tilde{G}$. $\square$

8 Kuratowski Subgraph Isolator

The non-planarity minors of Figure 6 can be used to find a Kuratowski subgraph in a non-planar graph (or a subgraph with at most $3n-5$ edges). Because the method is closely associated with the back edges, external face paths and DFS tree paths of the input graph, the linear time performance and correctness of the method are clear from the discussion.

The first step is to determine which non-planarity minor to use. Minors A to D are $K_{3,3}$ minors, so they may be used directly to find a subgraph homeomorphic to $K_{3,3}$. Minor E is a K_5 minor, so a few further tests are performed afterward to determine whether a subgraph homeomorphic to $K_{3,3}$ or K_5 can be obtained. To determine the minor, we first find an unembedded back edge (v, d), then search up the DFS tree path $T_{v,d}$ in $\tilde{G}$ to find the root v^c of a biconnected component on which the Walkdown failed. The Walkdown can then be reinvoked to determine whether the merge stack is empty (unless the merge stack is still available from the Walkdown that halted). Either way, the short-circuit edges should be deleted and $\tilde{G}$ should be properly oriented as described in Section 4 before proceeding.

If the merge stack is non-empty, then the desired biconnected component root r can be found at the top of the stack. Otherwise, we use v^c. The two external face paths from the selected root are searched for the stopping vertices x and y, then we search the lower external face path $(x, \ldots, y)$ for a pertinent vertex w that the Walkdown could not reach. Then, if the merge stack was non-empty, we invoke the minor A isolator (the isolators are described below).

If the merge stack is empty, then we must choose between minors B to E. If w has a pertinent externally active child biconnected component (check the last element of the pertinentRoots list), then we invoke the minor B isolator. Otherwise, we must find the highest x-y path by temporarily deleting the internal edges incident to v^c, then traversing the proper face bordered by v^c and its two remaining edges. Due to the removal of edges, the bounding cycle of the face will contain cut vertices, which can be easily recognized and eliminated as their cut vertices are visited for a second time during the walk. Once the x-y path is obtained, the internal edges incident to v^c are restored.

If either p_x or p_y is attached high, then we invoke the minor C isolator. Otherwise, we test for non-planarity minor D by scanning the internal vertices of the x-y path for a vertex z whose x-y path edges are not consecutive above the x-y path. If it exists, such a vertex z may be directly incident to v^c or it may have become a cut vertex during the x-y path test. Either way, we invoke the minor D isolator if z is found and the minor E isolator if not.

Each isolator marks the vertices and edges to be retained, then deletes all unmarked edges and merges biconnected component roots with their non-root counterparts. The edges are added and marked to complete the pertinent path from w to v and the external activity paths from x and y to ancestors of v. Minors B and E also require an additional edge to complete the external activity path for z. Finally, the tree path is added from v to the ancestor of least DFI associated with the external activity of x, y and (for minors B and E) z. Otherwise, we mark previously embedded edges along depth first search tree paths, the x-y path and v-z path, and the external face of the biconnected component containing the stopping vertices.

To exemplify marking an external activity path, we consider the one attached to x (in any of the non-planarity minors). If the least ancestor directly attached to x by a back edge (a value obtained during the lowpoint calculation) is less than v, then let u_x be that least ancestor, and let d_x be x. Otherwise, u_x is the lowpoint of the first child χ in the separatedDFSChildList of x, and let d_x be the neighbor of u_x in G with the least DFI greater than χ. We mark the DFS tree path from d_x to x and add and mark the edge (u_x, d_x). The external activity paths for y and, when needed, z are obtained in the same way.

Marking the pertinent path is similar, except that minor B requires the path to come from the pertinent child biconnected component containing z. In the other cases, the backedgeFlag flag tells whether we let d_w be w. If the backedgeFlag flag is clear or we have minor B, then we obtain the last element w^χ in the pertinentRoots list of w, then scan the adjacency list of v in G for the neighbor d_w with least DFI greater than χ. Finally, mark the DFS tree path d_w to w and add the edge (v, d_w).

To conclude the $K_{3,3}$ isolation for minor A, we mark the DFS tree path from v to the least of u_x and u_y and we mark the external face of the biconnected component rooted by r. For minor B, we mark the external face path of the biconnected component rooted by v^c and the DFS tree path from $\max(u_x, u_y, u_z)$ to $\min(u_x, u_y, u_z)$. The path from v to $\max(u_x, u_y, u_z)$, excluding endpoints, is not marked because the edge (u, v) in minor B is not needed to form a $K_{3,3}$. For the same reason, minors C and D omit parts of the external face of the biconnected component rooted by v^c, but both require the tree path v to $\min(u_x, u_y)$. Minor C omits the short path from p_x to v if p_x is attached high, and otherwise it omits the short path from p_y to v. Minor D omits the upper paths $(x, \ldots, v)$ and $(y, \ldots, v)$. In all cases, the endpoints of the omitted paths are not omitted.

Finally, the minor E isolator must decide between isolating a $K_{3,3}$ homeomorph and a K_5 homeomorph. Four simple tests are applied, the failure of which implies that minor E can be used to isolate a K_5 homeomorph based on the techniques described above. The first test to succeed implies the ability to apply the corresponding minor from Figure 7.

Minor E_1 occurs if the pertinent vertex w is not externally active (i.e. a second vertex z is externally active along the lower external face path strictly between p_x and p_y). If this condition fails, then $w = z$. Minor E_2 occurs if

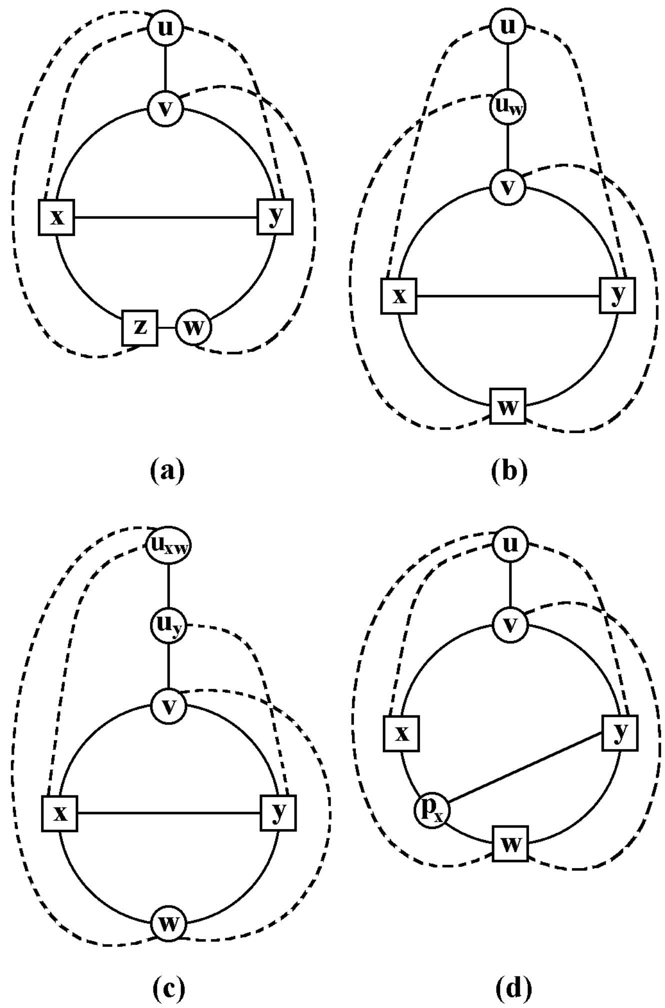

Figure 7: (a) Minor E_1, (b) Minor E_2, (c) Minor E_3, (d) Minor E_4

the external activity connection from w to an ancestor u_w of v is a descendant of u_x and u_y. Minor E_3 occurs if u_x and u_y are distinct and at least one is a descendant of u_w. Minor E_4 occurs if either $p_x \neq x$ or $p_y \neq y$.

As with minors A to D, there are symmetries to contend with and some edges that are not needed to form a $K_{3,3}$. For minors E_1 and E_2 it is easy to handle the symmetries because they reduce to minors C and A, respectively. Minor E_3 does not require (x, w) and (y, v) to form a $K_{3,3}$, and minor E$_4$ does not require (u, v) and (w, y) to form a $K_{3,3}$. Moreover, note that the omission of these edges must account for the fact that p_x or p_y may have been edge contracted into x or y in the depiction of the minor (e.g. eliminating (w, y) in minor E_4 corresponds to eliminating the path $(w, \ldots, p_y)$ but not $(p_y, \ldots, y)$).

As for symmetries, minor E_1 in Figure 7(a) depicts z between x and w along the path $(x, \ldots, z, \ldots, w, \ldots, y)$, but z may instead appear between w and y along the path $(x, \ldots, w, \ldots, z, \ldots, y)$. Also, Figure 7(c) depicts minor E_3 with u_x an ancestor of u_y, but u_y could instead be an ancestor of u_x. For minor E_4, Figure 7(d) depicts p_x distinct from x (and p_y can be equal to or distinct from y), but if $p_x = x$, then p_y must be distinct from y. Finally, note that the symmetric cases have different edges that have to be deleted to form a $K_{3,3}$.

9 Conclusion

This paper discussed the essential details of our new 'edge addition' planarity algorithm as well as a straightforward method for isolating Kuratowski subgraphs. With a graph theoretic basis, these algorithms simplify linear time graph planarity relative to prior approaches. An $O(n)$ implementation has been found to be competitive with implementations of the well-known prior methods, including being several times faster than LEDA implementations of the Hopcroft-Tarjan, PQ-tree and PC-tree methods [4, 20]. Our implementation has been rigorously tested on billions of randomly generated graphs and all connected graphs on 11 or fewer vertices (generated with McKay's nauty program [18]). Our implementation, as well as independent implementations such as [16], have required only a few weeks of implementation time.

Future research shall include extension of our method to outerplanarity, instances of the subgraph homeomorphism problem, the consecutive ones problem and interval graph recognition, and the generation of maximal planar subgraphs, visibility representations, and alternate planar embeddings. Our methods may also assist in the development of simplified, efficient embedders for the projective plane and the torus.

References

[1] K. S. Booth and G. S. Lueker. Testing for the consecutive ones property, interval graphs, and graph planarity using PQ–tree algorithms. *Journal of Computer and Systems Sciences*, 13:335–379, 1976.

[2] J. Boyer. *Simplified $O(n)$ algorithms for planar graph embedding, Kuratowski subgraph isolation and related problems.* Ph.D. Thesis, University of Victoria, 2001.

[3] J. Boyer and W. Myrvold. Stop minding your P's and Q's: A simplified $O(n)$ planar embedding algorithm. *Proceedings of the Tenth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 140–146, 1999.

[4] J. M. Boyer, P. F. Cortese, M. Patrignani, and G. Di Battista. Stop minding your P's and Q's: Implementing a fast and simple DFS-based planarity testing and embedding algorithm. In G. Liotta, editor, *Proceedings of the 11th International Symposium on Graph Drawing 2003*, Lecture Notes in Computer Science. Springer-Verlag, 2003.

[5] J. M. Boyer, C. G. Fernandes, A. Noma, and J. C. de Pina Jr. Correcting and implementing the PC-tree planarity algorithm. Submitted. Preprint at `http://www.pacificcoast.net/~lightning/CorrectPCTreePlanarity.ps`, August 2003.

[6] N. Chiba, T. Nishizeki, A. Abe, and T. Ozawa. A linear algorithm for embedding planar graphs using PQ–trees. *Journal of Computer and Systems Sciences*, 30:54–76, 1985.

[7] H. de Fraysseix and P. Rosenstiehl. A characterization of planar graphs by trémaux orders. *Combinatorica*, 5(2):127–135, 1985.

[8] N. Deo. Note on Hopcroft and Tarjan planarity algorithm. *Journal of the Association for Computing Machinery*, 23:74–75, 1976.

[9] G. Di Battista, P. Eades, R. Tamassia, and I. G. Tollis. *Graph Drawing: Algorithms for the Visualization of Graphs.* Prentice Hall, Upper Saddle River, NJ, 1999.

[10] S. Even and R. E. Tarjan. Computing an *st*-numbering. *Theoretical Computer Science*, 2:339–344, 1976.

[11] A. Gibbons. *Algorithmic Graph Theory.* Cambridge University Press, 1985.

[12] J. Hopcroft and R. Tarjan. Efficient planarity testing. *Journal of the Association for Computing Machinery*, 21(4):549–568, 1974.

[13] W.-L. Hsu and R. McConnell. PC-trees. In D. P. Mehta and S. Sahni, editors, *Handbook of Data Structures and Applications.* CRC Press, Boca Raton, FL, 2003 (in press).

[14] M. Jünger, S. Leipert, and P. Mutzel. Pitfalls of using PQ-trees in automatic graph drawing. In G. Di Battista, editor, *Proceedings of the 5th International Symposium on Graph Drawing '97*, volume 1353 of *Lecture Notes in Computer Science*, pages 193–204. Springer Verlag, Sept. 1997.

[15] A. Lempel, S. Even, and I. Cederbaum. An algorithm for planarity testing of graphs. In P. Rosenstiehl, editor, *Theory of Graphs*, pages 215–232, New York, 1967. Gordon and Breach.

[16] P. Lieby. Planar graphs. *The Magma Computational Algebra System*, http://magma.maths.usyd.edu.au/magma/htmlhelp/text1185.htm.

[17] R. J. Lipton and R. E. Tarjan. Applications of a planar separator theorem. *SIAM Journal of Computing*, 9(3):615–627, 1980.

[18] B. D. McKay. Practical graph isomorphism. *Congressus Numerantium*, 30:45–87, 1981.

[19] K. Mehlhorn and P. Mutzel. On the embedding phase of the Hopcroft and Tarjan planarity testing algorithm. *Algorithmica*, 16:233–242, 1996.

[20] A. Noma. Análise experimental de algoritmos de planaridade. Master's thesis, Universidade de São Paulo, May 2003. in Portuguese.

[21] E. M. Reingold, J. Nievergelt, and N. Deo. *Combinatorial Algorithms: Theory and Practice.* Prentice-Hall, Inc., Englewood Cliffs, New Jersey, 1977.

[22] W.-K. Shih and W.-L. Hsu. A simple test for planar graphs. In *Proceedings of the International Workshop on Discrete Mathematics and Algorithms*, pages 110–122, 1993.

[23] W.-K. Shih and W.-L. Hsu. A new planarity test. *Theoretical Computer Science*, 223:179–191, 1999.

[24] R. E. Tarjan. Depth-first search and linear graph algorithms. *SIAM Journal of Computing*, 1(2):146–160, 1972.

[25] R. Thomas. Planarity in linear time. `http://www.math.gatech.edu/~thomas/`, June 1997.

[26] D. B. West. *Introduction to Graph Theory*. Prentice Hall, Inc., Upper Saddle River, NJ, 1996.

[27] S. G. Williamson. Embedding graphs in the plane- algorithmic aspects. *Annals of Discrete Mathematics*, 6:349–384, 1980.

[28] S. G. Williamson. *Combinatorics for Computer Science*. Computer Science Press, Rockville, Maryland, 1985.