Connectivity of Graphs Under Edge Flips

Norbert Zeh

Technical Report CS-2003-07

September 15, 2003

Faculty of Computer Science
6050 University Ave., Halifax, Nova Scotia, B3H 1W5, Canada

Connectivity of Graphs Under Edge Flips

Norbert Zeh

Faculty of Computer Science, Dalhousie University,
6050 University Ave, Halifax, NS B3H 2Y5, Canada

nzeh@cs.dal.ca

Abstract

We study the following problem: Given a set V of n vertices and a set E of m edge pairs, we define a graph family $\mathcal{G}(V,E)$ as the set of graphs that have vertex set V and contain exactly one edge from each pair in E. We want to find a graph in $\mathcal{G}(V,E)$ that has the minimal number of connected components. We show that, if the edge pairs in E are non-disjoint, the problem is NP-hard. This is true even if an edge is not allowed to appear in more than two edge pairs and the union of all graphs in $\mathcal{G}(V,E)$ is planar. If the edge pairs are disjoint, we provide an $\mathcal{O}(n^2m)$-time algorithm that finds a graph in $\mathcal{G}(V,E)$ with the minimal number of connected components. Our proof of the latter statement is obtained using the concept of flipping edges in the graphs in $\mathcal{G}(V,E)$, where a flip in a graph $G \in \mathcal{G}(V,E)$ removes an edge $e \in G$ and replaces it with the other edge in the pair in E that contains e. We answer the following questions: Can every graph in $\mathcal{G}(V,E)$ be transformed into a graph in $\mathcal{G}(V,E)$ with the minimal number of connected components using a sequence of edge flips that steadily decrease the number of connected components in the graph? How long is the shortest such sequence (that is, how many flips are in this sequence)? Can we transform any graph in $\mathcal{G}(V,E)$ with at most k connected components into any other graph in $\mathcal{G}(V,E)$ with at most k connected components using a sequence of edge flips that never increase the number of connected components beyond k? How long does such a sequence have to be?

1 Introduction

Edge flips in abstract and geometric graphs are a well-studied concept. One restricts one's attention to a family $\mathcal{G}$ of graphs all of whose members have the same vertex set and the same number of edges. An edge flip is the operation of removing an edge e from the current graph G and replacing it with an edge $e' \neq e$ so that the resulting graph G' belongs to $\mathcal{G}$. Clearly, the requirement that $G' \in \mathcal{G}$ restricts the set of permissible edge flips because, by replacing an edge $e \in G$ with just any edge $e' \notin G$, one may produce a graph that does not belong to $\mathcal{G}$. The removal of an edge e may be altogether impermissible because there may not be any edge $e' \notin G$ so that $(G - e) \cup \{e'\}$ is in $\mathcal{G}$. Sometimes, the set of permissible edge flips is constrained even further using additional condition that disallow certain flips even though the resulting graph would be in $\mathcal{G}$. Classic examples of graph families $\mathcal{G}$ include spanning trees of a graph G and triangulations of a point set

P in the plane. In the former case, the replacement of an edge $e \in T$ with another edge $e' \notin T$ is permissible if $(T - e) \cup \{e'\}$ is a tree and $e' \in G$. In the latter case, we allow the replacement of edge e with edge e' if the resulting graph is a valid triangulation of P; that is, edge e' has to be a diagonal of the quadrilateral created by removing edge e, and edge e' must not intersect any edge in $G - e$. (Remember that G is a triangulation of a planar point set; that is, every vertex is augmented with a location, and every edge is represented as a straight line segment between the two points representing its endpoints.)

Usually, one studies questions such as: Can every graph in $\mathcal{G}$ be transformed into any other graph in $\mathcal{G}$ using a sequence of edge flips? How many edge flips are required? Or we may assign a quality measure to the members of $\mathcal{G}$, and we ask whether every member in $\mathcal{G}$ can be transformed into an optimal member of $\mathcal{G}$—optimal with respect to this quality measure—using a sequence of edge flips. If the answer is affirmative, one is again interested in the number of flips required to achieve this. Coming back to the example of spanning trees of a graph G, we may assign weights $\omega(e)$ to the edges of G and ask to find a minimum spanning tree of G with respect to these edge weights. In this case, we are interested in flips where $\omega(e) > \omega(e')$ because each such flip produces a spanning tree of weight less than the previous spanning tree. One can easily show that $n - 1$ edge flips suffice to transform any spanning tree into a minimum spanning tree. As for triangulations of planar point sets, the classical question is whether every triangulation of the point set can be transformed into the Delaunay triangulation [1] of the point set using a sequence of edge flips [3]. The flips we are interested in in this case are the so-called Delaunay flips, which take an edge that violates the Delaunay property and flip it. It is known than $\Omega(n^2)$ flips are necessary in the worst case [3] and $\mathcal{O}(n^2)$ flips suffice [3] to transform any triangulation into the Delaunay triangulation of the point set. For the case of arbitrary edge flips in triangulations of point sets and polygons in the plane, see for example [4, 6]. For a survey on edge flips in triangulated planar graphs, see [7].

In this paper, we study a very restricted class of graphs and a very restricted set of flips. The graph family $\mathcal{G} = \mathcal{G}(V, E)$ is defined by a vertex set V and a set E of edge pairs so that a graph belongs to $\mathcal{G}(V, E)$ if and only if it has vertex set V and contains exactly one edge from every pair in E. If the edge pairs are non-disjoint, it is possible that no graph in $\mathcal{G}(V, E)$ has a permissible flip; that is, for every graph $G \in \mathcal{G}(V, E)$, it is impossible to produce another graph in $\mathcal{G}(V, E)$ by removing a single edge from G and replacing it with another edge. If the edge pairs are disjoint, permissible flips are those that remove any edge e from the current graph G and replace it with the other edge in the edge pair containing e. We are interested in the connectivity of the graphs in $\mathcal{G}(V, E)$. The first question we ask is whether we can efficiently identify a graph in $\mathcal{G}(V, E)$ that has the minimal number of connected components among all graphs in $\mathcal{G}(V, E)$; we call this the *maximal connectivity problem* (MCP) ("maximal connectivity" because a graph with the minimal number of connected components is maximally connected in a sense). This question was the starting point of this research. It was raised by Edelsbrunner [2] as a graph-theoretic formulation of a problem arising in the repair of self-intersections of triangulated surfaces in space. For non-disjoint edge pairs, we prove that MCP is NP-hard, even if we restrict our attention to graph families $\mathcal{G}(V, E)$ so that the union of the graphs in $\mathcal{G}(V, E)$ is planar and no edge is allowed to be a member of more than two edge pairs in E; we call this version of the problem planar 2-MCP. More generally, we define k-MCP as MCP restricted to graph families $\mathcal{G}(V, E)$ where no edge is allowed to be in more than k pairs in E. For 1-MCP—the case of pairwise disjoint edge pairs—we provide an $\mathcal{O}(n^2m)$-time algorithm, where $n = |V|$ and $m = |E|$. We prove that every graph G in $\mathcal{G}(V, E)$ can be transformed into a graph with the minimal number of connected components using a sequence of at most m edge flips such that no flip in the sequence increases the number of connected components in the current graph. We do not know how to compute such a sequence efficiently; but we present an $\mathcal{O}(nm)$-time algorithm that computes a sequence of at most m flips such that no flip increases the number of connected components in the current graph and the final graph has less connected components than the initial graph G. In order to find a graph with the minimal number of connected components, we apply this algorithm at most $n - 2$ times, thereby obtaining the $\mathcal{O}(n^2m)$-time algorithm for finding a graph with the minimal number of connected components in $\mathcal{G}(V, E)$. Another question we ask is whether any graph $G \in \mathcal{G}(V, E)$ with at most k connected components can be transformed into any other graph in $\mathcal{G}(V, E)$ with at most k connected components using a sequence of flips that never produces a graph with more than k connected components. We prove that if k is greater than the minimal number of connected components of the graphs in $\mathcal{G}(V, E)$, then there exists such a sequence of at most $2m$ flips. If k is the minimal number of connected components of the graphs in $\mathcal{G}(V, E)$, we prove that such a sequence of flips may not exist.

The paper is organized as follows: In Section 2, we introduce the necessary notation and terminology. In Section 3, we provide the NP-hardness proof for planar k-MCP, for $k \geq 2$. In Sections 4 through 6, we study 1-MCP. In Section 4, we provide the proof that every graph $G \in \mathcal{G}(V, E)$ has a sequence of at most m flips that transforms it into a maximally connected graph in G and so that no flip increases the number of connected components in the current graph. In Section 5, we provide our $\mathcal{O}(nm)$-time algorithm for finding a sequence of at most m flips that transforms a given graph G into a graph with less connected components than G and never increases the number of connected components along the way. Using this algorithm, we obtain our $\mathcal{O}(n^2m)$-time algorithm for solving 1-MCP. In Section 6, we prove that, as long as there is a graph in $\mathcal{G}(V, E)$ with less than k connected components, every graph $G \in \mathcal{G}(V, E)$ with at most k connected components can be transformed into any other graph with at most k connected components using a sequence of at most $2m$ flips such that no flip produces a graph with more than k connected components. If k is the number of connected components of the maximal graphs in $\mathcal{G}(V, E)$, we prove that such a sequence may not exist. In Section 7, we conclude the discussion with a number of open problems.

2 Preliminaries

Given a set V of vertices and a set E of m edge pairs $P_1, \ldots, P_m$ over V, we define a graph family $\mathcal{G}(V, E)$ so that a graph G belongs to $\mathcal{G}(V, E)$ if and only if G has vertex set V and contains exactly one edge from every pair P_i, $1 \leq i \leq m$. We say that family $\mathcal{G}(V, E)$ is *k-thick* if every edge appears in at most k pairs in E. For a graph G, we define $\gamma(G)$ to be the number of connected components of G. For a family $\mathcal{G}(V, E)$, we define $\gamma(\mathcal{G}(V, E)) = \min\{\gamma(G) : G \in \mathcal{G}(V, E)\}$. We say that a graph $G \in \mathcal{G}(V, E)$ is *maximally connected* (or *maximal* for short) if $\gamma(G) = \gamma(\mathcal{G}(V, E))$. The *maximal connectivity problem* (MCP) is the problem of determining $\gamma(\mathcal{G}(V, E))$, for a given family $\mathcal{G}(V, E)$. MCP as a decision problem is to decide whether a given family $\mathcal{G}(V, E)$ contains a connected graph. We define k-MCP to be the maximal connectivity problem restricted to k-thick graph families $\mathcal{G}(V, E)$. Planar k-MCP is the k-MCP problem restricted to graph families $\mathcal{G}(V, E)$ such that the union of the graphs in $\mathcal{G}(V, E)$ is planar.

Given a 1-thick family $\mathcal{G}(V, E)$ with $E = \{P_1, \dots, P_m\}$ and an edge $e \in P_i$, we denote the other edge in P_i by $\bar{e}$ and call it the *complementary edge* of e. Given a graph $G \in \mathcal{G}(V, E)$ and an edge $e \in G$, the *flip* of edge e produces the graph $G\langle e \rangle = (G - e) \cup \{\bar{e}\}$; that is, we replace edge e with edge $\bar{e}$. This flip may affect the number of connected components in the graph. We call the flip *profitable* if $\gamma(G\langle e \rangle) \leq \gamma(G)$ and *augmenting* if $\gamma(G\langle e \rangle) < \gamma(G)$. We can extend this notion to sequences of flips: Let $e_1, \dots, e_q$ be a list of edges in G. In this paper, we will use $e_1, \dots, e_q$ to denote also the sequence of flips defined by flipping edges $e_1, \dots, e_q$ in this order. It will be clear from the context whether $e_1, \dots, e_q$ denotes a sequence of edge flips or a list of edges. We denote the graph obtained from a graph G by flipping edges $e_1, \dots, e_q$ as $G\langle e_1, \dots, e_q \rangle$. We say that the flip sequence $e_1, \dots, e_q$ is profitable if the flip of edge e_i is profitable for $G\langle e_1, \dots, e_{i-1} \rangle$, for all $1 \leq i \leq q$. A profitable flip sequence $e_1, \dots, e_q$ is augmenting if $\gamma(G\langle e_1, \dots, e_q \rangle) < \gamma(G)$.

Let e be an edge of a graph G in a 1-thick family $\mathcal{G}(V, E)$, let $V_1, \dots, V_k$ be the vertex sets of the connected components of G, and let $W_1, \dots, W_\ell$ be the vertex sets of the connected components of $G\langle e \rangle$. We say that the connected components of G are *invariant* under the flip of edge e if $k = \ell$, and there exists a permutation σ such that $V_i = W_{\sigma(i)}$, for all $1 \leq i \leq k$; that is, the vertex sets of the connected components stay invariant, while their edge sets do change.

3 Planar k-MCP is NP-Hard

Our proof that planar k-MCP is NP-hard for $k \geq 2$ uses a linear-time reduction from 3-SAT to planar 2-MCP. Let us first agree on the terminology we use to talk about formulas in 3-CNF. Given a Boolean variable x, we denote its negation as $\bar{x}$. A *literal* is a Boolean variable or its negation. A *clause* is the disjunction of literals: $C = \lambda_1 \vee \lambda_2 \vee \cdots \vee \lambda_k$. A Boolean formula F is in *conjunctive normal form* (CNF) if it is of the form $F = C_1 \wedge C_2 \wedge \cdots \wedge C_m$, where $C_1, \dots, C_m$ are clauses. Formula F is in 3-CNF if every clause C_i, $1 \leq i \leq m$, contains exactly three literals. In this case, we denote the literals in C_i as $\lambda_{i,1}$, $\lambda_{i,2}$, and $\lambda_{i,3}$. We denote the Boolean variables that occur in F as $x_1, \dots, x_n$. The satisfiability problem for formulas in 3-CNF (3-SAT) is defined as the problem of deciding whether a given formula in 3-CNF is satisfiable, that is, whether there exists a truth assignment to the variables $x_1, \dots, x_n$ so that F is true. It is well-known that 3-SAT is NP-complete [5]. Hence, if we can provide a polynomial-time reduction from 3-SAT to MCP, MCP is NP-hard. In this section, we provide such a reduction.

The central element used in a number of constructions in this paper is the "connector graph" shown in Figure 1. This graph is planar. Its edges are grouped into *disjoint* edge pairs as indicated by the numbering in Figure 1. We call this graph a connector graph because, no matter which edge we choose from each edge pair, the resulting graph is connected. We leave it as an exercise

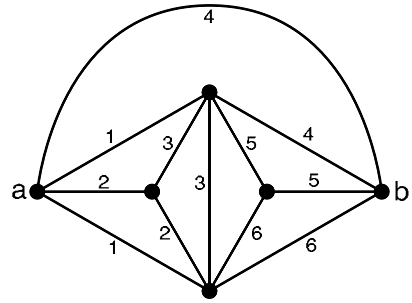

Figure 1
The connector graph. Two edges belong to the same edge pair if they have the same number.

to verify this fact. Since any subgraph of the connector graph that contains exactly one edge from every pair is connected, we can distinguish two of its vertices a and b—call them endpoints—and think of the graph as a "permanent" edge between vertices a and b; that is, this edge has to be present in every graph in the graph family. In all subsequent figures, we will represent such permanent edges as fat grey edges. In order to obtain a valid graph all of whose edges are assigned to edge pairs, each of these permanent edges has to be replaced by a connector graph containing the endpoints of the edge.

Given a formula F in 3-CNF with n variables $x_1, \ldots, x_n$ and m clauses $C_1, \ldots, C_m$, we construct a graph $G'(F)$ and assign its edges to appropriate edge pairs to obtain a graph family $\mathcal{G}(F)$. We make sure that no edge is in more than two pairs and that $G'(F)$, which is the union of the graphs in $\mathcal{G}(F)$, is planar. We construct family $\mathcal{G}(F)$ in such a way that it contains a connected graph if and only if F is satisfiable. Graph $G'(F)$ consists of a long chain of permanent edges (connector graphs sharing their endpoints); m vertices $\Gamma_1, \ldots, \Gamma_m$, one per clause in F; and a number of regular edges that represent the literals in the clauses of F (see Figure 2).

For every literal λ, let $\kappa(\lambda)$ be the number of clauses containing λ. For a variable x_i, we define $\kappa^*(x_i) = |\kappa(x_i) - \kappa(\bar{x}_i)|$. Let $\kappa^* = \sum_{i=1}^{n} \kappa(x_i)$. Then the chain of permanent edges contains $3m + 2\kappa^* + 2$ vertices. We denote the vertices in this chain by $\alpha_{\kappa^*}, \alpha_{\kappa^*-1}, \ldots, \alpha_1, \alpha', \beta', \beta_1, \beta_2, \ldots, \beta_{\kappa^*}$, $\gamma_{1,1}, \gamma_{1,2}, \gamma_{1,3}, \gamma_{2,1}, \gamma_{2,2}, \gamma_{2,3}, \ldots, \gamma_{m,1}, \gamma_{m,2}, \gamma_{m,3}$, in this order. Vertices $\Gamma_1, \ldots, \Gamma_m$ are not part of this chain. Besides the permanent edges, graph G' contains regular edges $e_{i,j} = (\Gamma_i, \gamma_{i,j})$, $1 \leq i \leq m$ and $1 \leq j \leq 3$, and f_k, $1 \leq k \leq \kappa^*$. We refer to an edge $e_{i,j}$ as a *literal edge* and to an edge f_k as a *dummy edge*.

Graph $G'(F)$ is obviously planar. The presence or absence of dummy edges does not affect the number of connected components in a subgraph of $G'(F)$. For such a subgraph to be connected, we need at least one edge $e_{i,j}$ per clause C_i to be present. Next we group literal edges and dummy edges into pairs so that, in every graph in $\mathcal{G}(F)$, a literal edge $e_{i,j}$ is present if and only if every edge $e_{i',j'}$ with $\lambda_{i,j} = \lambda_{i',j'}$ is present and every edge $e_{i'',j''}$ with $\lambda_{i,j} = \bar{\lambda}_{i'',j''}$ is absent. Hence, the presence and absence of edges in a graph in $\mathcal{G}(F)$ corresponds to a truth assignment to the variables $x_1, \ldots, x_n$; a graph G in $\mathcal{G}(F)$ is connected if and only if the corresponding truth assignment

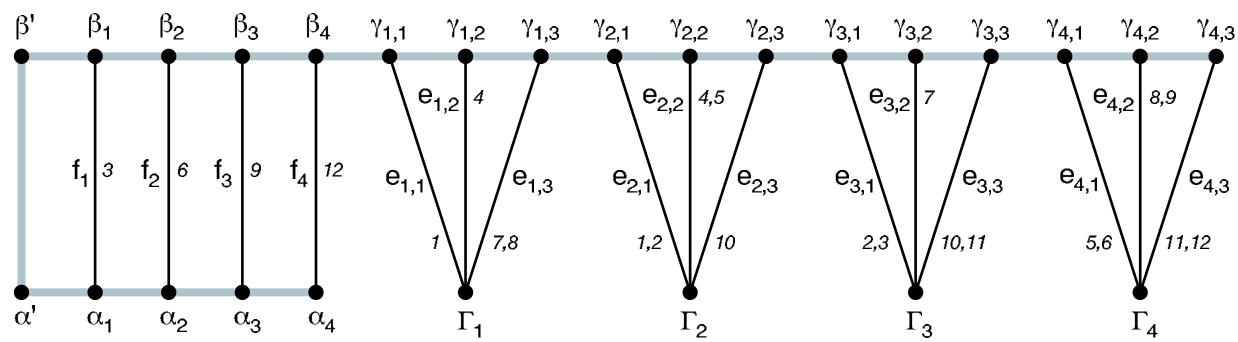

Figure 2

The graph derived from the formula $F = (x_1 \vee x_2 \vee \bar{x}_3) \wedge (\bar{x}_1 \vee \bar{x}_2 \vee x_4) \wedge (x_1 \vee x_3 \vee \bar{x}_4) \wedge (x_2 \vee x_3 \vee x_4)$. Regular edges are labelled with their names. The small italic labels identify the edge pairs that contain each edge.

satisfies F.

Consider a variable x_k and the literals $\lambda_{i_1,j_1}, \dots, \lambda_{i_q,j_q}$ so that $\lambda_{i_h,j_h} = x_k$ or $\lambda_{i_h,j_h} = \bar{x}_k$, for all $1 \le h \le q$. We assume w.l.o.g. that $\lambda_{i_1,j_1} = \cdots = \lambda_{i_r,j_r} = x_k$ and $\lambda_{i_{r+1},j_{r+1}}, \dots, \lambda_{i_q,j_q} = \bar{x}_k$, for some $1 \le r \le q$. We also assume that $\kappa^*(x_k) = \kappa(x_k) - \kappa(\bar{x}_k)$, that is, there are at least as many positive literals x_k in F as negative literals $\bar{x}_k$. Then we choose a set of $s = \kappa^*(x_k)$ dummy edges $f_{l_1}, \dots, f_{l_s}$ that have not been included in any pairs yet. We define the following edge pairs, where $t = q - r$: $\{e_{i_h,j_h}, e_{i_{h+r},j_{h+r}}\}$, for $1 \le h \le t$; $\{e_{i_{h+r},j_{h+r}}, e_{i_{h+1},j_{h+1}}\}$, for $1 \le h \le \max(t, r-1)$; $\{e_{i_{h+t},j_{h+t}}, f_{l_h}\}$, for $1 \le h \le s$; and $\{f_{l_h}, e_{i_{h+t+1},j_{h+t+1}}\}$, for $1 \le h < s$. Intuitively, this looks like in Figure 3; that is, we construct a "path" of edge pairs, where edges corresponding to positive and negative literals alternate until we run out of negative literals; once this happens, we place a dummy edge between every pair of consecutive positive literals. Clearly, this ensures that all edges corresponding to literal x_k are either all present or all absent and that the former is the case if and only if all edges corresponding to literal $\bar{x}_k$ are absent, and the latter is the case if and only if all edges corresponding to literal $\bar{x}_k$ are present. Note that, by creating κ^* dummy edges, we made sure that we have enough dummy edges to complete this construction for all variables $x_1, \dots, x_n$ while using every dummy edge in the creation of edge pairs for exactly one variable x_k. This ensures that, indeed, every edge is in at most two edge pairs, as illustrated in Figure 2.

Lemma 1 *There is a connected graph in $\mathcal{G}(F)$ if and only if F is satisfiable.*

Proof. It follows from the above discussion that every graph $G \in \mathcal{G}(F)$ corresponds to a valid truth assignment to variables $x_1, \dots, x_k$. If G is connected, there is at least one edge connecting each clause node Γ_i to one of its corresponding nodes $\gamma_{i,j}$. The presence of this edge corresponds to the presence of a true literal in clause C_i, namely $\lambda_{i,j}$. Hence, formula F is satisfied by the truth assignment corresponding to G. Conversely, if F is satisfiable, there exists a truth assignment so

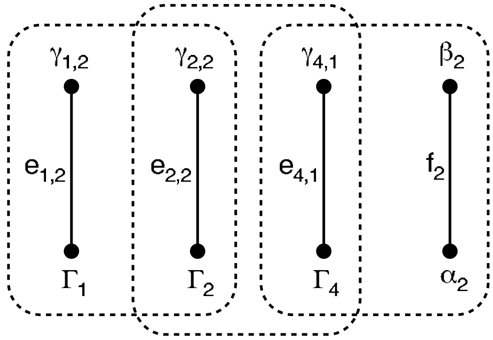

Figure 3
The "path" of alternating literal edges for the variable x_2 in the formula $F = (x_1 \vee x_2 \vee \bar{x}_3) \wedge (\bar{x}_1 \vee \bar{x}_2 \vee x_4) \wedge (x_1 \vee x_3 \vee \bar{x}_4) \wedge (x_2 \vee x_3 \vee x_4)$. The membership in edge pairs is shown by dotted boxes around every such pair.

that every clause contains at least one true literal. Hence, in the graph $G \in \mathcal{G}(F)$ corresponding to this truth assignment, every clause node Γ_i has at least one incident edge, and G is connected. $\square$

From Lemma 1, we obtain our first main result.

Theorem 1 *Planar k-MCP is NP-hard, for any $k \geq 2$.*

Proof. Graph $G'(F)$ has size $\Theta(m)$ and is planar. It is easy to see that every edge is contained in at most two edge pairs. Hence, graph family $\mathcal{G}(F)$ is 2-thick, and there are $\Theta(m)$ edge pairs. The construction of graph $G'(F)$ and of the edge pairs from a given formula F can be carried out in $\mathcal{O}(m)$ time. Hence, if planar k-MCP is solvable in polynomial time, for any $k \geq 2$, so is 3-SAT. This shows that planar k-MCP is NP-hard. $\square$

Corollary 1 *k-MCP is NP-hard, for any $k \geq 2$. In particular, MCP is NP-hard.*

4 Existence of Profitable and Augmenting Flips and Sequences

Given the NP-hardness of k-MCP for $k \geq 2$, we restrict our attention to 1-MCP in the remainder of this paper. We will apply a "greedy" strategy to solve 1-MCP: Given a 1-thick graph family $\mathcal{G}(V,E)$, we start with an arbitrary graph $G \in \mathcal{G}(V,E)$ and then compute an augmenting sequence of flips that transforms G into a maximal graph. In this section, we prove that such a sequence always exists; in fact, there always is such a sequence consisting of at most m flips. In the next section, we show how to find such a sequence consisting of at most $(n-2)m$ flips in $\mathcal{O}(n^2m)$ time. Efficiently finding an augmenting flip sequence for a non-maximal graph G that transforms G into a maximal graph and consists of at most m flips is an open problem.

Let us start slowly, by showing that every non-maximal graph has a profitable, though not necessarily augmenting flip. We prove that there is such a flip that leaves the connected components of G invariant unless it is augmenting.

Lemma 2 *Every non-maximal graph G in a family $\mathcal{G}(V, E)$ has an augmenting flip or a profitable flip so that the connected components of G are invariant under this flip.*

Proof. Observe that if G is a forest, it must be maximal. Hence, G is not a forest. Then we can choose an edge e that is on a cycle in G. Removing edge e from G does not alter the connected components of G. Adding edge $\bar{e}$ decreases the number of connected components by one if the endpoints of $\bar{e}$ are in different connected components of G; the flip is augmenting in this case. If the endpoints of $\bar{e}$ are in the same connected component of G, the addition of edge $\bar{e}$ leaves the connected components of G invariant. □

Given a graph G, we call a flip of an edge $e \in G$ *greedy* if the endpoints of edge $\bar{e}$ are in different connected components of G. A sequence $e_1, \ldots, e_q$ of edge flips is greedy if, for every $1 \leq i \leq q$, the flip of edge e_i is greedy for $G\langle e_1, \ldots, e_{i-1}\rangle$. The following two observations suggest that restricting our attention to greedy flips is a good idea.

Observation 1 *A flip is augmenting for a graph G if and only if it is greedy and removes an edge from a cycle in G.*

Observation 2 *An edge $e \in G$ so that the flip of edge e is greedy, but not augmenting, for G is a cut-edge of G.*

Another way to interpret Observation 2 is that non-augmenting greedy flips leave all cycles in G intact. This will be important in a number of arguments we make. Also observe that for every greedy flip e, edge $\bar{e}$ is a cut-edge of $G\langle e\rangle$. Armed with these two observations, we are now ready to prove that not every non-maximal graph has an augmenting flip.

Lemma 3 *Not every non-maximal graph has an augmenting flip.*

Proof. Consider Figure 4a. By Observation 1, every augmenting flip has to remove an edge in the bold triangle. However, none of these flips is augmenting, as each of them inserts an edge between two vertices in the same connected component. Still, graph G is non-maximal because the sequence $2, 5$ transforms G into the graph in Figure 4b, which has one less connected component than G. □

Even though, by Lemma 3, not every non-maximal graph has an augmenting flip, every such graph has an augmenting sequence of flips.

Lemma 4 *Every non-maximal graph has an augmenting sequence of flips.*

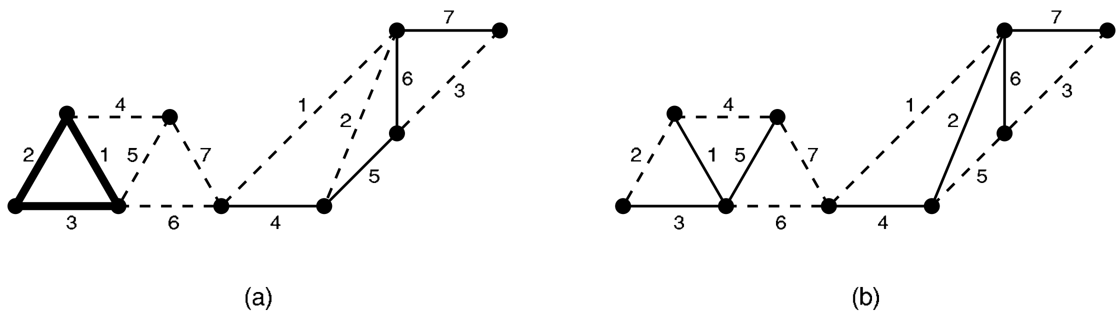

Figure 4
(a) A graph G with no augmenting flip. Solid edges are in the graph; dotted edges are not in the graph and can be exchanged for edges in the graph through flips. (b) A graph in the same family with one less connected component than G.

Proof. Let G be a non-maximal graph in a 1-thick family $\mathcal{G}(V, E)$, let $k = \gamma(G)$ be the number of connected components of G, and let $G' \in \mathcal{G}(V, E)$ be a graph with less than k connected components. Graph G' exists because $\gamma(G) > \gamma(\mathcal{G}(V, E))$. We prove the lemma by induction on the number of edges in G that are not in G'. If there is exactly one such edge e, flipping this edge transforms G into G'; this one-flip sequence is augmenting. So assume that G contains $r > 1$ edges that are not in G' and that the lemma holds for every graph G'' that contains less than r edges that are not in G'. Since G' has one less connected component than G, there has to be an edge e in G' that connects two vertices in different connected components of G. The flip of edge $\bar{e}$ is greedy for G; that is, this flip is profitable, if not augmenting for G. If $G\langle\bar{e}\rangle$ has less than k connected components, then the one-flip sequence $\bar{e}$ is augmenting for G. Otherwise, $G\langle\bar{e}\rangle$ is a graph with k connected components that contains less than r edges that are not in G'. Hence, by the inductive hypothesis, there exists an augmenting sequence $e_1, \ldots, e_t$ of flips for $G\langle\bar{e}\rangle$. The sequence $\bar{e}, e_1, \ldots, e_t$ is augmenting for G. □

The following corollary is an immediate consequence of the proof of Lemma 4.

Corollary 2 *Every non-maximal graph in $\mathcal{G}(V, E)$ has an augmenting sequence of at most m flips.*

In fact, the arguments in the proof of Lemma 4 are easily generalized to show the following lemma. All we have to do is choose the graph G' in the proof to be maximal and, instead of stopping the construction as soon as we find a graph with less than k connected components, continue flipping edges until we obtain a maximal graph.

Lemma 5 *Every non-maximal graph G in $\mathcal{G}(V, E)$ has an augmenting sequence of at most m flips that transforms G into a maximal graph.*

5 Finding Augmenting Sequences of Flips

Given that every graph can be transformed into a maximal graph using a sequence of at most m flips, each of which is profitable, we would like to be able to compute such a sequence. Finding an efficient algorithm that computes such a sequence consisting of at most m flips is an open problem. However, in this section, we provide an algorithm that finds an augmenting flip sequence of length at most m for any non-maximal graph. Applying this procedure at most $n-2$ times, we obtain a maximal graph. Hence, by concatenating the computed flip sequences, we obtain a sequence of at most $(n-2)m$ flips that transforms the given graph into a maximal one.

Given a graph $G \in \mathcal{G}(V, E)$, we use an auxiliary directed graph H, which is derived from G, to find an augmenting sequence of flips for G. We prove that every shortest path between certain vertices in H corresponds to an augmenting sequence of flips for G and that at least one such path exists if G is non-maximal. Hence, all we have to do is apply breadth-first search to H to either find such a shortest path and report the corresponding sequence of flips or output that G is maximal if no such path exists.

The vertex set of H is the edge set of G. For two edges e and f of G, there is an edge (e, f) (directed from e to f) in H if the endpoints of $\bar{e}$ are in the same connected component of G, f is a cut-edge of G, and f is on every path in G connecting the two endpoints of $\bar{e}$. We allow edges (e, e) if edge e is itself a cut-edge and is on every path connecting the two endpoints of $\bar{e}$. We call a source e of H a *root* if it does not correspond to a cut-edge of G; that is, there exists a simple cycle in G that contains e. We call a sink e a *leaf* if edge $\bar{e}$ has its endpoints in different connected components of G. To be more explicit, a vertex e may be a source if edge e is not a cut-edge or it is a cut-edge that is not on any path connecting the endpoints of an edge $\bar{f}$. Only in the former case is the vertex considered to be a root. A vertex e may be a sink because the complementary edge $\bar{e}$ of e has its endpoints in different connected components or because $\bar{e}$ connects two vertices in a cycle in G. Only in the former case is the vertex considered to be a leaf. Every augmenting sequence of flips has to flip an edge corresponding to a root, because it has to break at least one cycle in G. It also has to flip at least one edge corresponding to a leaf in H because, otherwise, the flips in the sequence cannot reduce the number of connected components. In particular, the augmenting sequence constructed in the proof of Lemma 4 starts with a greedy flip, which corresponds to a leaf in H, and ends with an augmenting flip, which, by Observation 1, corresponds to a root.

Our goal now is to show that graph H contains a root-to-leaf path if G is not maximal and that every shortest such path corresponds to an augmenting sequence of flips. So assume that G is not maximal, and let $e_1, \dots, e_q$ be an augmenting sequence of greedy flips for G. As shown in the proof of Lemma 4, such a sequence exists if G is not maximal. Moreover, assume that $e_1, \dots, e_q$ is the shortest augmenting sequence of greedy flips for G. The following two lemmas are the first steps toward showing that there exists a root-to-leaf path in H.

Lemma 6 *Edges $e_1, \dots, e_{q-1}$ are cut-edges of G.*

Proof. Assume the contrary and choose j minimal so that e_j is not a cut-edge; that is, $e_1, \dots, e_{j-1}$ are cut-edges. We claim that $e_1, \dots, e_j$ is an augmenting greedy sequence of flips. This would contradict the assumption that $e_1, \ldots, e_q$ is the shortest greedy sequence of flips and would hence prove the lemma. Sequence $e_1, \ldots, e_j$ is certainly greedy, as every subsequence $e_1, \ldots, e_h$ of a greedy sequence of flips must itself be greedy. To see that sequence $e_1, \ldots, e_j$ is augmenting, observe that graph $G\langle e_1, \ldots, e_{j-1}\rangle$ has no less connected components than G, because sequence $e_1, \ldots, e_q$ is augmenting; edge $\bar{e}_j$ has its endpoints in different connected components of graph $G\langle e_1, \ldots, e_{j-1}\rangle$; and the cycles of G are invariant under deletion of cut-edges from G. Hence, e_j is not a cut-edge of $G\langle e_1, \ldots, e_{j-1}\rangle$. To summarize, flipping edge e_j removes an edge from a cycle and inserts an edge between two connected components of $G\langle e_1, \ldots, e_{j-1}\rangle$. Therefore, by Observation 1, the flip of edge e_j is augmenting for $G\langle e_1, \ldots, e_{j-1}\rangle$, and $\gamma(G\langle e_1, \ldots, e_j\rangle) < \gamma(G\langle e_1, \ldots, e_{j-1}\rangle) \leq \gamma(G)$. □

Lemma 7 *Edge e_q is not a cut-edge of G.*

Proof. Since sequence $e_1, \ldots, e_q$ is a shortest augmenting greedy sequence of flips, no subsequence $e_1, \ldots, e_j$, $j < q$, is augmenting. Hence, the flip of edge e_q is augmenting for $G\langle e_1, \ldots, e_{q-1}\rangle$. By Observation 1, this implies that edge e_q is not a cut-edge of $G\langle e_1, \ldots, e_{q-1}\rangle$. If e_q is a cut-edge of G, we choose j minimal so that e_q is not a cut-edge of $G\langle e_1, \ldots, e_j\rangle$. Since e_q is a cut-edge of $G\langle e_1, \ldots, e_{j-1}\rangle$, the insertion of edge $\bar{e}_j$ must create a cycle in $G\langle e_1, \ldots, e_j\rangle$. But then the endpoints of $\bar{e}_j$ are in the same connected component of $G\langle e_1, \ldots, e_{j-1}\rangle$, contradicting the greediness of sequence $e_1, \ldots, e_q$. Thus, edge e_q is not a cut-edge of G. □

Lemma 7 implies that e_q is a root in H. Since the endpoints of edge e_1 are in different connected components of G, by the greediness of sequence $e_1, \ldots, e_q$, edge e_1 is a leaf of H. Hence, to prove that graph H contains a root-to-leaf path if G is not maximal, it suffices to show that there exists a path from e_q to e_1. We are unable to show exactly this; but we can show that if there is no path from e_q to e_1, there exists a path from e_q to another leaf.

Lemma 8 *If G is not maximal, then there exists a root-to-leaf path in H.*

Proof. Consider the same augmenting greedy flip sequence $e_1, \ldots, e_q$ as before. (For this proof, it is irrelevant whether this sequence is shortest possible.) We have just observed that e_1 is a leaf, and e_q is a root of H. We show that for every edge e_i, $1 \leq i \leq q$, there exists a path from e_i to a leaf in H. Hence, this is true in particular for e_q, which is a root. The proof is by induction. Since e_1 is itself a leaf, the claim holds for e_1. So assume that $i > 1$ and that the claim holds for $e_1, \ldots, e_{i-1}$. If e_i is a leaf, the claim holds for e_i. Otherwise, the endpoints of $\bar{e}_i$ are in the same connected component of G. But they are in different connected components of $G\langle e_1, \ldots, e_{i-1}\rangle$; so there must be an edge e_j, $j < i$, that is on all paths in G connecting the endpoints of $\bar{e}_i$. This implies that e_j is an out-neighbour of e_i in H. By the induction hypothesis, there exists a path from e_j to a leaf. This implies that there exists a path from e_i to a leaf. □

Given that graph H is guaranteed to contain a root-to-leaf path if G is not maximal, it is natural to ask what the relationship between root-to-leaf paths in H and augmenting flip sequences for G is. Our first question is whether just any root-to-leaf path corresponds to an augmenting sequence

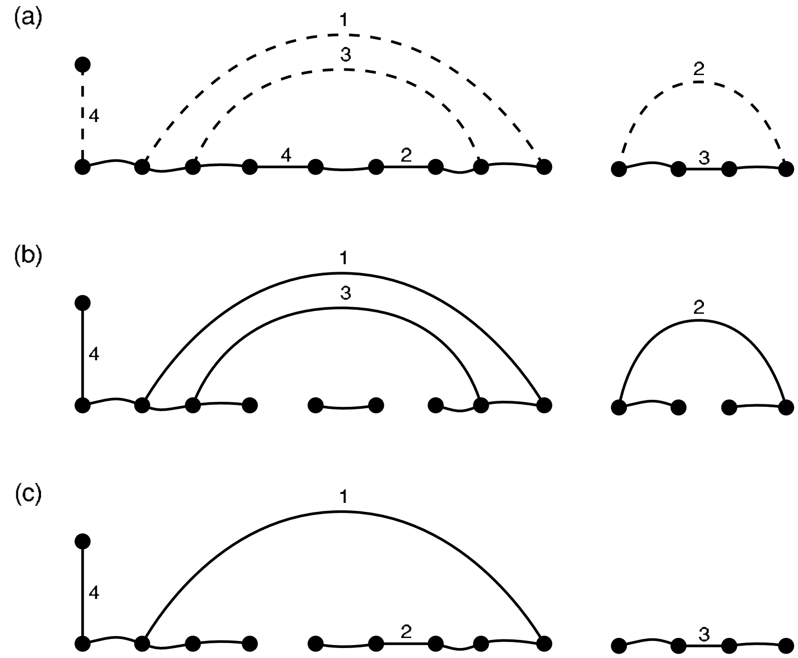

Figure 5
(a) A portion of a graph G and four edge pairs. Solid edges are in G. Dashed edges are not in G. The edge in G that belongs to pair 1 is not shown; it belongs to a cycle in G. (b) The graph $G\langle e_1, \ldots, e_4\rangle$ has the same number of connected components as graph G. (c) The graph $G\langle e_1, e_4\rangle$ has one less connected components than G.

of flips. The example in Figure 5 provides a negative answer. However, the problem with this example seems to be that we did not choose the shortest path from the root e_1 to the leaf e_4. The shortest path is (e_1, e_4); as shown in Figure 5c, the sequence e_1, e_4 is augmenting. Next we prove that every shortest root-to-leaf path in H corresponds to an augmenting sequence of flips. To prove this fact, we make use of the following two results.

Lemma 9 *For a shortest root-to-leaf path $(e_1, \ldots, e_q)$ in H and any $1 \leq i < q$, edges $\bar{e}_i$ and e_{i+1} are in a common cycle of graph $G\langle e_1, \ldots, e_i\rangle$.*

Proof. Assume that the lemma does not hold. Then there exists an i so that edges $\bar{e}_i$ and e_{i+1} are not in a common cycle in $G\langle e_1, \ldots, e_i\rangle$. Choose i minimally so. Since edge (e_i, e_{i+1}) exists in H, edges $\bar{e}_i$ and e_{i+1} are in a common cycle in $G \cup \{\bar{e}_i\}$. Choose j maximal so that $\bar{e}_i$ and e_{i+1} are in a common cycle of $G\langle e_1, \ldots, e_h\rangle \cup \{\bar{e}_i\}$, for all $0 \leq h \leq j$. Then $j < i$, and edges $\bar{e}_i$ and e_{i+1} do not belong to a common cycle in the graph $G\langle e_1, \ldots, e_{j+1}\rangle \cup \{\bar{e}_i\}$. Hence, edge e_{j+1} is on the cycle in $G\langle e_1, \ldots, e_j\rangle \cup \{\bar{e}_i\}$ containing edges e_{i+1} and $\bar{e}_i$ (see Figure 6a). Since $\bar{e}_j$ and e_{j+1} belong to the same cycle in $G\langle e_1, \ldots, e_j\rangle$, the removal of edge e_{j+1} still leaves a cycle containing $\bar{e}_i$ and e_{i+1}

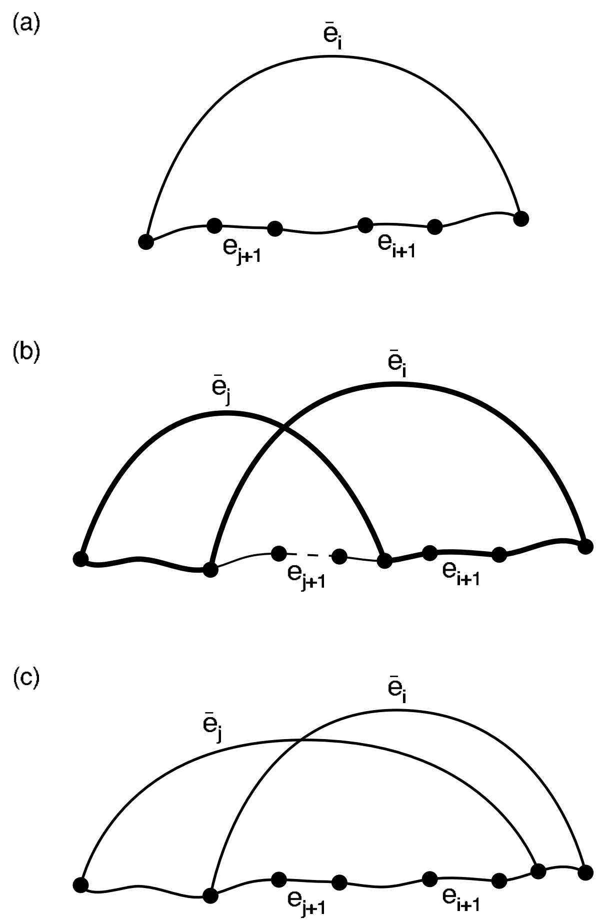

Figure 6
(a) Edges e_{i+1} and e_{j+1} belong to a cycle in $G\langle e_1, \ldots, e_j\rangle \cup \{\bar{e}_i\}$. (b) If edge e_{i+1} is not on the cycle in $G\langle e_1, \ldots, e_j\rangle$ containing edges e_{j+1} and $\bar{e}_j$, then there exists a cycle (bold) in $G\langle e_1, \ldots, e_{j+1}\rangle \cup \{\bar{e}_i\}$ that contains edges $\bar{e}_i$ and e_{i+1}. (c) The case when edge e_{i+1} is on the cycle in $G\langle e_1, \ldots, e_j\rangle$ containing edges e_{j+1} and $\bar{e}_j$.

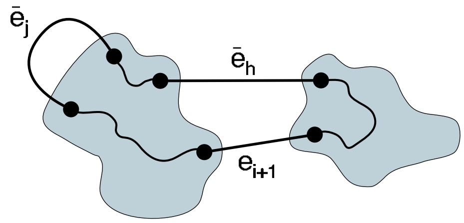

Figure 7
The proof that there must be an edge $(\bar{e}_h, e_{i+1})$ in H, where $h < i$, if edge e_{i+1} is on the cycle in $G\langle e_1, \ldots, e_j \rangle$ that contains edges e_{j+1} and $\bar{e}_j$.

(Figure 6b), unless the cycle containing $\bar{e}_j$ and e_{j+1} also contains e_{i+1} (Figure 6c). In the former case, we obtain a contradiction to the assumption that no cycle containing $\bar{e}_i$ and e_{i+1} exists in $G\langle e_1, \ldots, e_{j+1} \rangle \cup \{\bar{e}_i\}$. We prove that in the latter case, there exists a shorter path from e_1 to e_q in H, which contradicts the assumption that path $e_1, \ldots, e_q$ is a shortest root-to-leaf path in H. This proves that edges $\bar{e}_i$ and e_{i+1} do belong to a common cycle in $G\langle e_1, \ldots, e_i \rangle$.

Since edge (e_j, e_{j+1}) exists in H, edge e_{j+1} is a cut-edge on the path in G connecting the two endpoints of $\bar{e}_j$. If e_{i+1} is also on this path, then edge (e_j, e_{i+1}) exists, and $(e_1, \ldots, e_j, e_{i+1}, \ldots, e_q)$ is a shorter path from e_1 to e_q in H, a contradiction. So assume that e_{i+1} is not on this path. Edge e_{i+1} is a cut-edge of G, and initially both endpoints of edge $\bar{e}_j$ are in the same connected component of $G - e_{i+1}$. Moreover, none of edges $\bar{e}_1, \ldots, \bar{e}_j$ connects two vertices in different connected components of G. Hence, the only way to create a path in $G\langle e_1, \ldots, e_{j-1} \rangle$ that connects the endpoints of edge $\bar{e}_j$ and contains edge e_{i+1} is the addition of an edge $\bar{e}_h$, $h < j$, whose endpoints are in different connected components of $G - e_{i+1}$, but in the same connected component of G (see Figure 7). Then, however, edge (e_h, e_{i+1}) exists in H, and the path $(e_1, \ldots, e_h, e_{i+1}, \ldots, e_q)$ is a shorter path from e_1 to e_q in H, again a contradiction. □

Corollary 3 *Let $e_1, \ldots, e_q$ be a shortest root-to-leaf path in H; let $V_1, \ldots, V_k$ be the vertex sets of the connected components of G; and let $W_1, \ldots, W_l$ be the vertex sets of the connected components of $G\langle e_1, \ldots, e_j \rangle$, for an arbitrary $1 \leq j < q$. Then $k = l$, and there exists a permutation σ so that $V_i = W_{\sigma(i)}$, for all $1 \leq i \leq k$.*

Using Lemma 9 and Corollary 3, we can now prove that every shortest root-to-leaf path in H corresponds to an augmenting sequence of flips for G.

Lemma 10 *A shortest root-to-leaf path in H corresponds to an augmenting sequence of flips for G.*

Proof. Consider a shortest root-to-leaf path $e_1, \ldots, e_q$ in H. By Corollary 3, all flips in the sequence $e_1, \ldots, e_{q-1}$ are profitable, but not augmenting. By Lemma 9, edge e_q belongs to a cycle in $G\langle e_1, \ldots, e_{q-1} \rangle$; by Corollary 3, edge $\bar{e}_q$ has its endpoints in different connected components of

$G\langle e_1, \ldots, e_{q-1}\rangle$, because this is true in G. Hence, by Observation 1, the flip of edge e_q is augmenting for $G\langle e_1, \ldots, e_{q-1}\rangle$, and the whole sequence $e_1, \ldots, e_q$ is augmenting for G. $\square$

Given this correspondence between root-to-leaf paths in H and augmenting sequences for G, we can find an augmenting sequence of flips for G in $\mathcal{O}(nm)$ time: First we create the vertex set of graph H by adding a vertex for every edge of G. Next we identify the cut-edges of G and label all those vertices in H as roots whose corresponding edges in G are not cut-edges. We contract every 2-edge connected component into a single vertex. Call the resulting graph G'. We compute its connected components, which are trees, and root each such tree at an arbitrary vertex. To identify the edge set of H and the leaves of H, we scan the set of complementary edges of the edges in G. We can discard edges which have become loops as the result of the contraction of the 2-edge connected components of G, as they run parallel to a path in a 2-edge connected component of G and hence neither have any out-neighbours nor are leaves. We mark a vertex e in H as a leaf if the endpoints of edge $\bar{e}$ are in different connected components of G'. For every edge $\bar{e}$ that is not a loop and that has its endpoints in the same connected component of G', we add an edge (e, f), for every edge f on the tree path connecting the endpoints of $\bar{e}$. These edges can be identified by traversing paths from the endpoints of $\bar{e}$ to their LCA in the tree. Since there are at most $n-1$ cut-edges in G, the edge set of G' has size at most $n-1$; the vertex set has size at most n. There are m edges $\bar{e}$. Hence, after constructing G', which takes $\mathcal{O}(n+m)$ time, it takes $\mathcal{O}(nm)$ time to construct H. Now we decide whether there exists a root-to-leaf path in H and, if so, find a shortest such path by running BFS simultaneously from all roots. That is, we place all the roots at the first level of the BFS and then grow the BFS-forest as usual level by level. This takes $\mathcal{O}(nm)$ time. If G is not maximal, our discussion implies that this procedure finds a root-to-leaf path. Since we use BFS to find it, it is a shortest such path. Hence, we report the sequence of flips corresponding to the vertices on the path as an augmenting sequence of flips.

Theorem 2 *It takes $\mathcal{O}(nm)$ time to decide whether a graph $G \in \mathcal{G}(V, E)$ is maximal and, if not, find an augmenting sequence of at most m flips.*

Corollary 4 *It takes $\mathcal{O}(n^2 m)$ time to compute a maximal graph $G \in \mathcal{G}(V, E)$. The procedure to compute G performs at most $(n-2)m$ edge flips, each of which is profitable.*

Proof. We start with an arbitrary graph $G \in \mathcal{G}(V, E)$, which can be constructed in $\mathcal{O}(n+m)$ time. If there is at least one edge pair in E, graph G has at most $n-1$ connected components; otherwise, G is maximal, and we are done. A maximal graph in $\mathcal{G}(V, E)$ has at least one connected component. Hence, we can construct a sequence $G = G_1, \ldots, G_t$ of at most $n-2$ graphs so that G_t is maximal and G_i is obtained from G_{i-1} through an augmenting sequence of flips. Such a sequence can be found in $\mathcal{O}(nm)$ time, for every graph G_{i-1}, and we repeat this computation at most $n-2$ times. Hence, we obtain G_t from G in $\mathcal{O}(n^2 m)$ time. Every graph G_i is obtained from graph G_{i-1} using an augmenting sequence of at most m edge flips. Hence, we perform at most $(n-2)m$ flips in total, and every flip is profitable. $\square$

6 Connectivity of Sub-Families Under Edge Flips

Since every graph in a 1-thick family $\mathcal{G}(V,E)$ can be transformed into a maximal graph, an interesting question to ask is whether for any two graphs G_1 and G_2 in $\mathcal{G}(V,E)$ with at most k connected components, there exists a flip sequence $e_1,\ldots,e_q$ that transforms G_1 into G_2 and such that $\gamma(G_1\langle e_1,\ldots,e_i\rangle) \leq k$, for all $1 \leq i \leq q$. We call such a sequence *semi-profitable*, as not all its flips are necessarily profitable, but the sequence maintains a certain minimal connectivity throughout. Another question to ask is how many flips such a sequence has to contain. Formally, this translates into the following problem: We consider a given family $\mathcal{G}(V,E)$ to be a graph whose nodes correspond to the graphs in $\mathcal{G}(V,E)$. There is an edge between two nodes in this graph if the two corresponding graphs can be transformed into each other by flipping a single edge. For any $k \geq \gamma(\mathcal{G}(V,E))$, we denote by $\mathcal{G}_{\leq k}(V,E)$ the subgraph of $\mathcal{G}(V,E)$ induced by all nodes whose corresponding graphs have at most k connected components. Then we ask whether $\mathcal{G}_{\leq k}(V,E)$ is connected, for any $k \geq \gamma(\mathcal{G}(V,E))$, and, if so, what is its diameter. We prove that $\mathcal{G}_{\leq k}(V,E)$ is connected, for every $k > \gamma(G(V,E))$, and that $\mathcal{G}_{\gamma(\mathcal{G}(V,E))}(V,E)$ may be disconnected. We also show that, for $k > \gamma(\mathcal{G}(V,E))$, $\mathcal{G}_{\leq k}(V,E)$ has diameter at most $2m$. To show this, we prove a fact that is in a sense orthogonal to Lemma 5. While Lemma 5 shows that every non-maximal graph can be transformed into *some* maximal graph using an *augmenting* sequence of at most m flips, we prove next that every non-maximal graph can be transformed into *any* maximal graph using a *semi-profitable* sequence of at most m flips, which is not necessarily augmenting.

Lemma 11 *For any graph family $\mathcal{G}(V,E)$ and any $k > \gamma(\mathcal{G}(V,E))$, $\mathcal{G}_{\leq k}(V,E)$ is connected and has diameter at most $2m$.*

Proof. The proof is an adaptation of the proof of Lemma 4. We prove that, for any two graphs G_1 and G_2 in $\mathcal{G}_{\leq k}(V,E)$, where G_2 is maximal, there exists a semi-profitable sequence $e_1,\ldots,e_q$ of at most m flips that transforms G_1 into G_2. Note that this implies the lemma because, for any two graphs G_1 and G_2 in $\mathcal{G}_{\leq k}(V,E)$, we can choose a maximal graph G_3 and two semi-profitable sequences sequences $e_1,\ldots,e_q$ and $e'_1,\ldots,e'_r$ that transform G_1 and G_2 into G_3 and such that $q,r \leq m$. Then the sequence $e_1,\ldots,e_q,\bar{e}'_r,\ldots,\bar{e}'_1$, which has length $q+r \leq 2m$, transforms G_1 into G_2 and is semi-profitable for G_1.

So let $G_1 \in \mathcal{G}_{\leq k}(V,E)$, let G_2 be maximal, and let $e_1,\ldots,e_\ell$ be the edges in G_1 that are not in G_2. We prove by induction on ℓ that there exists a permutation $\sigma : \{1,\ldots,\ell\} \to \{1,\ldots,\ell\}$ such that the sequence $e_{\sigma(1)},\ldots,e_{\sigma(\ell)}$ is semi-profitable for G_1. Clearly, $G_1\langle e_{\sigma(1)},\ldots,e_{\sigma(\ell)}\rangle = G_2$. Since $\ell \leq m$, this proves our claim.

For $\ell = 1$, there is exactly one edge $e_1 \in G_1$ that is not in G_2. Flipping this edge produces G_2; that is, $G_1\langle e_1\rangle = G_2$. Since we flip exactly one edge and both G_1 and G_2 have no more than k connected components, this establishes the base case of our inductive proof.

So assume that $\ell > 1$ and that the claim holds for any two graphs G'_1 and G'_2 in $\mathcal{G}_{\leq k}(V,E)$ such that G'_2 is maximal and G'_1 contains less than ℓ edges that are not in G'_2. If G_1 is maximal, we can choose any edge in G_1 that is not in G_2, say e_1, and flip it. Observe that graph $G_1\langle e_1\rangle$ is in $\mathcal{G}_{\leq k}(V,E)$. This is true because G_1 is maximal and $k > \gamma(\mathcal{G}(V,E))$. Moreover, the edges in graph $G_1\langle e_1\rangle$ that

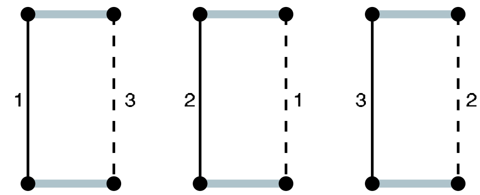

Figure 8
A family $\mathcal{G}(V, E)$ of graphs so that $\mathcal{G}_{\leq\gamma(\mathcal{G}(V,E))}(V, E)$ is disconnected.

are not in G_2 are edges $e_2, \ldots, e_\ell$. By the induction hypothesis, we can find a permutation of these edges so that the resulting sequence is semi-profitable for $G_1\langle e_1\rangle$. The sequence of flips starting with edge e_1, followed by this semi-profitable sequence of flips for $G_1\langle e_1\rangle$, is semi-profitable for G_1.

If G_1 is non-maximal, there has to be an edge $\bar{e}_i \in G_2$ that has its endpoints in two different connected components of G_1. As argued before, the flip of edge e_i is profitable for G_1 in this case, so that $G_1\langle e_i\rangle$ is in $\mathcal{G}_{\leq k}(V, E)$. The edges in $G_1\langle e_i\rangle$ that are not in G_2 are edges $e_1, \ldots, e_{i-1}, e_{i+1}, \ldots, e_\ell$. By the inductive hypothesis, we can find a permutation of these edges so that the resulting sequence is semi-profitable for $G_1\langle e_i\rangle$. The sequence of flips starting with edge e_i, followed by this semi-profitable sequence of flips for $G_1\langle e_i\rangle$, is semi-profitable for G_1. This completes the inductive step and thus finishes the proof of the lemma. □

The statement of Lemma 11 is true for any $k > \gamma(\mathcal{G}(V, E))$. Our proof of the lemma exploits the fact that, for a maximal graph, any edge flip produces a graph that is in $\mathcal{G}_{\leq k}(V, E)$. Is this just a caveat of the proof, or can $\mathcal{G}_{\gamma(\mathcal{G}(V,E))}(V, E)$ indeed be disconnected? The next lemma proves that the latter is the case.

Lemma 12 *There exists a graph family $\mathcal{G}(V, E)$ so that $\mathcal{G}_{\leq\gamma(\mathcal{G}(V,E))}(V, E)$ is disconnected.*

Proof. Consider the graph shown in Figure 8. The horizontal edges are permanent; that is, they are to be replaced by connectors again. We can construct two graphs G_1 and G_2 by including the permanent edges and the solid vertical edges in G_1 and including the permanent edges and the dashed vertical edges in G_2. Observe that, starting with G_1, flipping any edge in a connector does not bring us any closer to G_2; flipping any vertical edge increases the number of connected components by one. Hence, there is no semi-profitable sequence of flips that transforms G_1 into G_2. □

7 Open Problems

The algorithm for finding an augmenting sequence of flips takes $\mathcal{O}(nm)$ time. We do believe that there should be an $\mathcal{O}(n + m)$ time algorithm for this problem. If this is true, we would be able to compute a maximal graph in a 1-thick graph family $\mathcal{G}(V,E)$ in $\mathcal{O}(n(n+m))$ time. What is the lower bound for the running time of any algorithm that solves 1-MCP. Can one obtain a more efficient algorithm for planar 1-MCP? Can we improve the output computed by the algorithm; more precisely, is there an algorithm that computes an augmenting sequence $e_1, \dots, e_q$ of at most m flips for any graph $G \in \mathcal{G}(V,E)$ so that $G\langle e_1, \dots, e_q \rangle$ is maximal? By Lemma 5, such a sequence exists; but we do not know how to compute it efficiently. Also note that, if we drop the requirement that every flip in the sequence be profitable, at most $n-1$ flips suffice to obtain a maximal graph from any graph in $\mathcal{G}(V,E)$. (It suffices to make sure that the edges in a spanning forest of a maximal graph are present.) Is there an efficient algorithm to find these flips?

As for the diameter of $\mathcal{G}_{\leq k}(V,E)$, for $k > \gamma(\mathcal{G}(V,E))$, we believe that it is at most am, for some constant $a < 2$, but at least $m + b$, for some constant $b \geq 1$. Determining the correct bound is an open problem.

Finally, one can ask similar questions about the behaviour of the biconnected components of a graph under edge flips. We could define the maximal biconnectivity problem (MBP) as the problem of finding a graph in a family $\mathcal{G}(V,E)$ that has the minimal number of biconnected components. We can then ask whether (planar) k-MBP is NP-hard for $k \geq 2$. We believe that the answer to this question is "yes" and that some adaptation of our proof from Section 3 proves this. The more interesting question is how to adapt our algorithm for 1-MCP to obtain an efficient solution for 1-MBP.

Acknowledgments. I would like to thank Anil Maheshwari for some insightful discussion on the subject and Ferran Hurtado for giving a talk that made me see 1-MCP as an edge-flip problem.

References

1. B. Delaunay. Sur la sphère vide. *Izvestiya Akademii Nauk SSSR, Otdelenie Matematicheskii i Estestvennyka Nauk*, 7:793–800, 1934.

2. H. Edelsbrunner. Personal communication. 2003.

3. S. Fortune. Voronoi diagrams and Delaunay triangulations. In *Computing in Euclidean Geometry*, D. Z. Hu and F. K. Wang, eds., pages 225–265. World Scientific, Singapore, 2nd edition, 1995.

4. J. Galtier, F. Hurtado, M. Noy, S. Pérennes, and J. Urrutia. Simultaneous edge flipping in triangulations. *International Jounal on Computational Geometry and Applications*, 13(2):113–133, 2003.

5. M. R. Garey and D. S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W. H. Freeman, San Francisco, 1979.

6. F. Hurtado, M. Noy, and J. Urrutia. Flipping edges in triangulations. *Discrete and Computational Geometry*, 22:333–346, 1999.

7. S. Negami. Diagonal flips of triangulations on surfaces. *Yokohama Mathematical Journal*, 47:1–40, 1999.