Electronic Colloquium on Computational Complexity, Report No. 50 (2009)

Logspace reduction of directed reachability for bounded genus graphs to the planar case*

Jan Kynčl and Tomáš Vyskočil

Department of Applied Mathematics and Institute for Theoretical Computer Science (ITI)[1],
Charles University, Faculty of Mathematics and Physics,
Malostranské nám. 25, 118 00 Praha 1, Czech Republic;
kyncl@kam.mff.cuni.cz, whisky@kam.mff.cuni.cz

Abstract

Directed reachability (or briefly reachability) is the following decision problem: given a directed graph G and two of its vertices s, t, determine whether there is a directed path from s to t in G. Directed reachability is a standard complete problem for the complexity class NL. Planar reachability is an important restricted version of the reachability problem, where the input graph is planar. Planar reachability is hard for L and is contained in NL but is not known to be NL-complete or contained in L. Allender et al. showed that reachability for graphs embedded on the torus is logspace-reducible to the planar case. We generalize this result to graphs embedded on a fixed surface of arbitrary genus.

1 Introduction

Directed reachability (or briefly reachability) is a standard complete problem for the complexity class NL (nondeterministic logspace). The problem is defined as follows: given a directed graph G and two of its vertices s, t, determine whether there is a directed path from s to t in G. Such a path is briefly called a *directed s-t-path*. In our definition of a directed graph we allow two edges in opposite directions between the same pair of vertices, but we do not allow loops. In the literature the reachability problem is also referred to as (directed) s-t-connectivity. Planar reachability is an important restricted version of the reachability problem,

*The results of this paper were obtained during the DIMACS/DIMATIA REU program at Rutgers University in June and July 2007, supported by the project ME 886 of the Ministry of Education of the Czech Republic.

[1]ITI is supported by the project 1M0545 of the Ministry of Education of the Czech Republic.

ISSN 1433-8092

where the input graph is planar. It can be assumed that the graph is given with its planar embedding since a planar embedding of a planar graph can be constructed in logspace [2]. Planar reachability is hard for L and is contained in NL but is not known to be NL-complete or contained in L. Recently Bourke, Tewari and Vinodchandran [3] improved the complexity upper bound by showing that planar reachability is in UL (unambiguous logspace), which is the class of decision problems that can be solved by a non-deterministic log-space Turing machine that has at most one accepting computation on any input. Allender et al. [1] showed that reachability for graphs embedded on the torus is logspace-reducible to planar reachability, and hence belongs to UL as well. We generalize this result to graphs embedded on arbitrary fixed surface, orientable or non-orientable.

Theorem 1. *For each fixed connected compact surface S, the reachability problem for graphs embedded in S is logspace-reducible to planar reachability.*

Note that we have to assume that the graph is given together with its embedding in S, since it is not known whether such an embedding can be found in logspace if S has positive genus.

Thierauf and Wagner [7] generalized the result of Bourke, Tewari and Vinodchandran [3] in another direction by showing that reachability in $K_{3,3}$-free graphs and K_5-free graphs is logspace-reducible to planar reachability.

For some particular classes of graphs the reachability problem is known to be in L. The most important examples include undirected graphs [6], (directed) series parallel graphs [4], and several subclasses of grid graphs [1].

2 Preliminaries

Let S be an (orientable or non-orientable) connected compact surface of genus $g > 0$. Let Π be an embedding of a given directed graph G with n vertices in S. If S is an orientable surface, the embedding is given by the rotation system of G. The *rotation system* is a set of rotations of all vertices, where the *rotation* of a vertex v is the (clockwise) cyclic permutation of edges incident with v. If S is non-orientable, the embedding is given by the rotation system of G and an orientation function $\lambda : E(G) \to \{1, -1\}$ which assigns -1 to an edge e if and only if e changes orientation. We may assume that G is connected since the reachability problem for undirected graphs is in L [6]. We may also assume that the surface S is minimal for Π, that is, each face of the embedding is homeomorphic to a disc. Such an embedding is called a 2-*cell embedding*.

The *Euler characteristic* of S, denoted by $\chi(S)$, is defined as $\chi(S) = 2 - 2g$ if S is orientable, and $\chi(S) = 2 - g$ if S is non-orientable. Equivalently, if v, e and f denote the number of vertices, edges and faces of a 2-cell embedding of G in S, then $\chi(S) = v - e + f$. In this way we can also define the Euler characteristic $\chi(\Pi)$ of the embedding Π. The Euler characteristic of Π can be computed in logspace in the following way. We enumerate all the faces of Π by traversing along facial walks (in clockwise direction), starting from every ordered pair of adjacent vertices, which we call a *vector*. We label all the vectors as $u_1, u_2, \ldots, u_{2m}$. In the beginning we set $f = 0$. In step i ($i = 1, 2, \ldots, 2m-1$), we list all facial walks starting with vectors $u_1, u_2, \ldots, u_i$ and check whether the vector u_{i+1} appeared in the list, as a part of some walk. If not, the facial walk starting with u_{i+1} determines a new face and we increase f by one.

3 Proof of Theorem 1

3.1 Main idea

Let s and t be two given vertices of G. We describe a logspace construction of a planar graph G'' containing vertices s'' and t'' such that G'' contains a directed s''-t''-path if and only if G contains a directed s-t-path.

We follow the approach of Allender et al. [1]. The main idea is to find cycles of G that do not separate S, and cut the surface S along them. This operation reduces the genus, so after finitely many steps we get a planar embedding of some resulting graph G'. The cutting operation, however, can destroy the connectivity properties of the graph. This is fixed by gluing several copies of G' together. If the original surface S has negative Euler characteristic, a naive gluing can produce a graph of exponential size, which is not constructible in logspace. We fix this problem by showing that a graph of polynomial size is sufficient to restore the connectivity properties; this is essentially the main new ingredient in this proof. The core idea is that there are only polynomially many distinct "topological" types of (directed) paths in G, as each path in G is a non-crossing curve in S.

3.2 Finding a non-separating cycle

The construction starts with finding a spanning tree T of G [2, 5]. The construction of the spanning tree reduces to undirected reachability which is in L due to Reingold [6]. The spanning tree is not stored in the memory, however. Instead we get a function which takes an edge e as an input and answers TRUE if and only if $e \in T$, using only a logarithmic amount of memory. This is, in fact, a common interpretation of any "logspace construction" (of objects of polynomial size).

A cycle in G is called *non-separating* if the corresponding closed curve in the embedding does not separate S. The following lemma shows that a non-separating cycle can be found efficiently. It is stated in [1] for orientable surfaces only, but the proof works for non-orientable surfaces as well.

Lemma 2. [1] *Let G be a connected graph 2-cell embedded in a surface S of positive genus. Let T be a spanning tree of G. Then there exists an edge $e \in E(G)\setminus T$ such that the (fundamental) cycle contained in $T \cup \{e\}$ is non-separating.*

A cycle C in G is *one-sided* if C has an odd number of orientation-changing edges (those with $\lambda(e) = -1$), and *two-sided* otherwise. Note that on an orientable surface every cycle is two-sided, and any one-sided cycle (on a non-orientable surface) is non-separating. Given a cycle C in G, we can test in logspace whether C is one-sided, by traversing along the cycle and multiplying the orientations of its edges. In case C is two-sided, the test whether C is non-separating can be also performed in logspace [1], by checking for an existence of a vertex v that is connected by a path (internally disjoint with C) to each side of C. It follows by Lemma 2 that a non-separating cycle can be found in logspace.

3.3 Cutting operation

Now we describe the cutting operation. When cutting G along a two-sided cycle $C = v_1 v_2 \dots v_k$, we replace the cycle C with two new cycles $C' = v_1' v_2' \dots v_k'$ and $C'' = v_1'' v_2'' \dots v_k''$; see Figure 1. For each edge $v_i w$ ($w \in G - C$) on the left side of C, we create an edge $v_i' w$ and for each edge $v_i w$ on the right side of C we create an edge $v_i'' w$. The new edges are directed in the same way as the corresponding edges in G. The rotation system of the new graph is determined by the rotation system of G in the obvious way; the main difference is that the copies of adjacent edges of C become adjacent in the rotations of the vertices v_i' and v_i''. The orientation of each new edge is the same as the orientation of the corresponding original edge. Alternatively, we can perform several switching operations to make the orientations of all the the cycle edges positive. A *switching operation* at a vertex v reverses the rotation of v and changes the orientations of all its incident edges. The switching operation does not change the embedding as it preserves all facial cycles.

The result of the cutting operation is an embedding into a surface of Euler characteristic $\chi(S)+2$. This surface is obtained by patching the two holes created by the cutting with two discs, so the cycles C' and C'' become facial cycles of the new embedding.

When cutting of G along a one-sided cycle $C = v_1 v_2 \dots v_k$, we replace C with a new cycle $D = v_1' v_2' \dots v_k' v_1'' v_2'' \dots v_k''$; see Figure 2. Although D is one-sided, we can still distinguish the left and the right side of the path $P = v_1 v_2 \dots v_k$. For each edge $v_i w$ ($w \in G - C$) on the left side of P we create an edge $v_i' w$ and for each edge $v_i w$ on the right side of P we create an edge $v_i'' w$. The directions and the orientations of new edges and the rotation system of the new embedding are determined similarly as in the previous case. The result is an embedding into a surface of Euler characteristic $\chi(S) + 1$. This surface is obtained by patching the hole created by the cutting with a disc. The cycle D becomes a facial cycle of

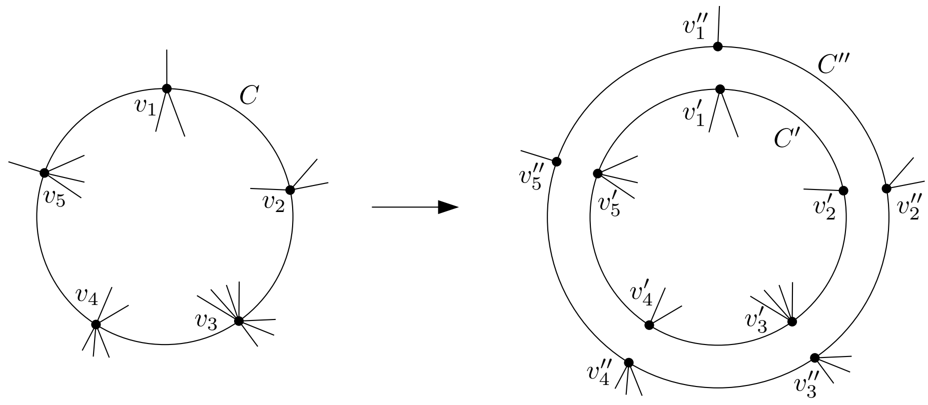

Figure 1: Cutting along two-sided cycle C. For simplicity, the orientations of the edges of C are chosen to be positive and the directions of the edges are not drawn.

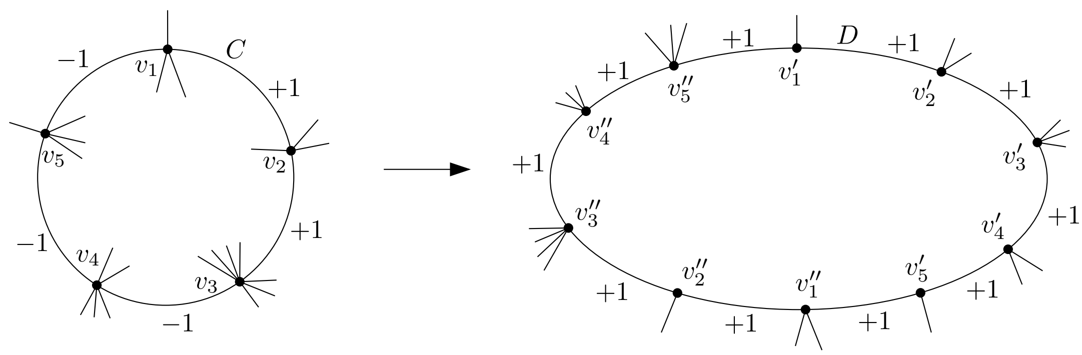

Figure 2: Cutting along one-sided cycle C on a non-orientable surface. The labels $+1$ and -1 denote the orientations of the cycle edges. For simplicity, the directions of the edges are not drawn.

the new embedding.

3.4 Reducing the genus

Starting from the given embedding Π of the graph G, we sequentially cut the graph along non-separating cycles, as long as the Euler characteristic of the embedding is smaller than 2. Since each cutting operation increases the Euler characteristic (and decreases the genus), after at most g cuttings we get a planar embedding Π' of a graph G'. Since the cutting operation preserves facial cycles, the embedding Π' contains facial cycles $C_1', C_1'', C_2', C_2'', \ldots, C_{g'}', C_{g'}''$ and $D_1, D_2, \ldots, D_{g''}$, where each pair C_i', C_i'' corresponds to a cutting along a two-sided cycle C_i and each D_i corresponds to a cutting along a one-sided cycle. The faces bounded by the cycles C_i', C_i'' and D_i are called *holes*. Each cutting increases the number of vertices at most twice and can be performed in logspace. Since at most g cuttings are performed, the graph G' has at most $2^g \cdot n$ vertices and can be constructed in logspace as well.

The vertices s and t might be split into more vertices during the cuttings. In such case we just choose one of the copies of s and one of the copies of t and call them s' and t', respectively. The existence of a directed s'-t'-path in G' implies the existence of a directed s-t-path in G, but not the other way. It might happen that a directed path in G was cut into several pieces during the cuttings.

3.5 Restoring connectivity

To restore the connectivity of G', we can glue a certain number of copies of the embedding Π' together. Two copies can be glued by the cycle C_i' in one copy with the cycle C_i'' in the other copy so that corresponding vertices and edges are identified. Or we can glue them by the cycle $D_i = v_1'v_2' \dots v_k'v_1''v_2'' \dots v_k''$ in one copy with the same cycle D_i in the other copy, so that the vertices v_i', v_i'' from one copy are identified with v_i'', v_i' from the other copy, respectively.

From the description of the cutting operations it follows that the total number of holes in Π' is equal to $2 - \chi(S) \leq 2g$. In case S is the projective plane, the embedding Π' has a single hole D_1, so we can glue at most two copies of Π' together to get an embedding in the sphere (which is the universal cover of the projective plane). In all other cases $\chi(S) \leq 0$ so Π' has at least two holes and we can glue arbitrary number of copies together. In case S is the torus or the Klein bottle ($\chi(S) = 0$), we have 2 holes in Π' so the gluing will result in an embedding in the cylindrical surface. For all S with negative Euler characteristic we have at least 3 holes in Π' and the gluing will result in an embedding in a tree-like surface (homeomorphic to a subset of the plane).

We start with Π_0', a "root" copy of Π', and glue a copy of Π' to each hole of Π_0'. Then in $2^g \cdot n - 1$ subsequent steps we glue a copy of Π' to each hole of the embedding constructed in the previous step; see Figure 1. We obtain a planar embedding of a graph H that has better connectivity than G'. More precisely, G has a directed s-t-path if and only if s' in Π_0' is connected by a directed path to one of the copies of t' in H. This follows from the observation that each directed s-t-path P in G that was cut into i components (including one-vertex components that arise when more cutting cycles pass through the same vertex) is restored after at most $i - 1$ steps, as a *lifted* directed path $\widetilde{P}$ in H.

In case S is the torus, the projective plane or the Klein bottle, the graph H has a linear or a quadratic size and can be constructed in logspace. In all other cases, however, H has exponential size and thus cannot be constructed in logspace.

We show that instead of H we can take a subgraph of H of polynomial size with the same connectivity.

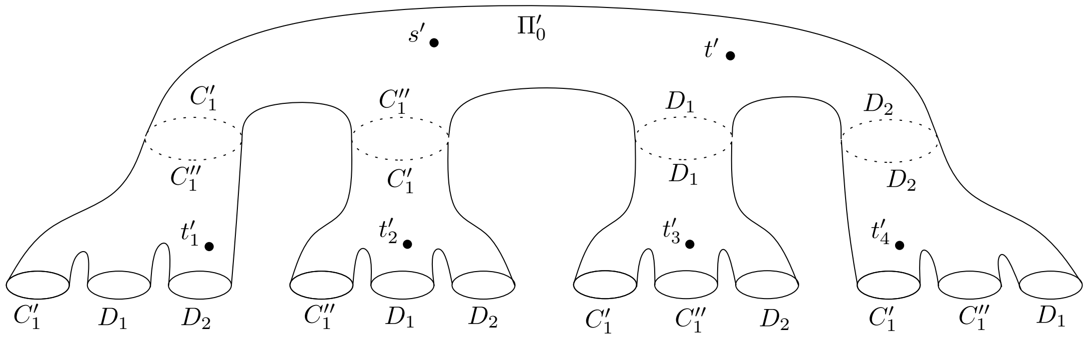

Figure 3: The surface obtained after the first step of gluing copies of Π'.

Lemma 3. *There is a logspace-constructible directed plane graph K (of polynomial size) containing vertices $s', t'_1, t'_2, \ldots, t'_N$ such that s is connected by a directed path to t in G if and only if s' is connected by a directed path to one of the vertices t'_i in K.*

For the final step of the reduction we use the following result of Allender et al. [1].

Lemma 4. [1] *Given a directed planar graph H containing vertices $s', t'_1, t'_2, \ldots,$ t'_N, there is a logspace-constructible directed planar graph H' containing vertices s'' and t'' such that H' contains a directed s''-t''-path if and only if H contains a directed path from s' to one of the vertices t'_i.*

The proof of Lemma 4 uses a reduction of planar reachability to a special case where both vertices s and t are on the outer face. Such reduction is performed for each pair s', t'_i in H, yielding a planar graph H_i with s' and t'_i on its outer face. Then a new vertex s'' is added and connected by a directed edge to a copy of s' in each H_i. Similarly a new vertex t'' is added and each vertex t'_i is connected by a directed edge to t''.

3.6 Proof of Lemma 3

The main idea of the construction is that only polynomially many branches of the tree structure of H are needed to cover the lifts of all s-t-paths in G. We show how to identify these branches and how to enumerate them in logspace.

Let P be a (not necessarily directed) s-v-path in G (where v is an arbitrary vertex). Let $(P_1, P_2, \ldots, P_k)$, $k \leq 2^g \cdot n$, denote the sequence of paths in G' that arise from P by the cuttings. Each P_i is a curve (possibly with zero length) with endpoints on the boundary of some hole (except P_1 starts in s' and P_k ends in v', where v' is one of the copies of v in G').

The *type* of P is a sequence $(X_1, X_2, \ldots, X_{k-1})$ of cycles C'_i, C''_i or D_i such that X_i contains the terminal vertex of P_i and corresponds to the cutting operation

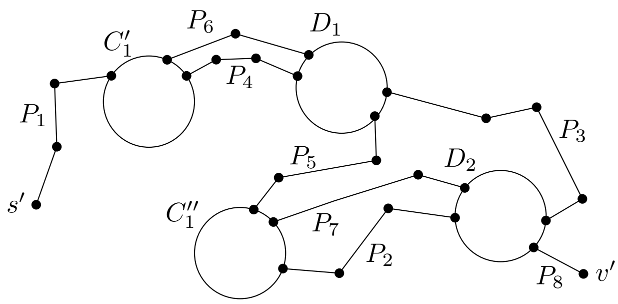

Figure 4: A path P of type $(C_1', D_2, D_1, C_1', D_1, C_1', D_2)$ cut into eight pieces.

that separated the components P_i and P_{i+1}. See Figure 4. If X_i is the cycle C_j' (C_j'', D_j), then let X_i' denote the cycle C_j'' (C_j', D_j, respectively).

The type of P determines through which copies of G' in H (we shall call these copies the *regions* of H) the lifted path $\widetilde{P}$ passes. The number of different types of directed paths in G may be exponential, but we show that all the corresponding lifted paths are contained in a union of only polynomial number of regions. To show this it is enough to prove that there are only polynomially many regions containing the terminal vertex of some lifted s-v-path. Every such region is called the *terminal region* of the corresponding path.

If $X_i' = X_{i+1}$ for some $i \leq k-2$, then the path P_{i+1} starts and ends on the boundary of the same hole, which means that the the paths P_i and P_{i+2} terminate in the same region. Thus after removing X_i and X_{i+1} from the sequence $X = (X_1, X_2, \dots, X_{k-1})$ we get a sequence Y which is a type of a path that we get from P by replacing P_{i+1} and parts of P_i and P_{i+2} in a neighborhood of X_i with a curve drawn along the cycle X_i; see Figure 5. Hence Y determines the same terminal region as X. By a slight abuse of notation, we still call the resulting curve an s-v-path, although the replacement curve does not consist of the edges of G. However, it still passes through each vertex of G at most once. Even if more reductions occur at the same cycle X_i, we can make sure that the new path does not cross itself. An s-v-path P in G and its type $(X_1, X_2, \dots, X_{k-1})$ are called *reduced* if there is no i such that $X_i' = X_{i+1}$. From the preceding observation it follows that for each s-v-path P in G there is a reduced s-v-path with the same terminal region.

Let P be a reduced s-v-path in G. We define an undirected *characteristic multigraph* of P, which we denote by $M(P)$, together with its planar embedding; see Figure 6. The vertex set of $M(P)$ consists of the points s' and v' in the embedding Π' and points c_i', c_i'' and d_i chosen in the interiors of the holes bounded by the cycles C_i', C_i'' and D_i, respectively. For each path P_i we draw an edge between the vertices x_{i-1}' and x_i inside the cycles X_{i-1}' and X_i (except for P_1 and P_k, where one of the vertices is s' or v', respectively). The edge is drawn along the curve representing the path P_i in Π' and extended to the points x_{i-1}' and x_i

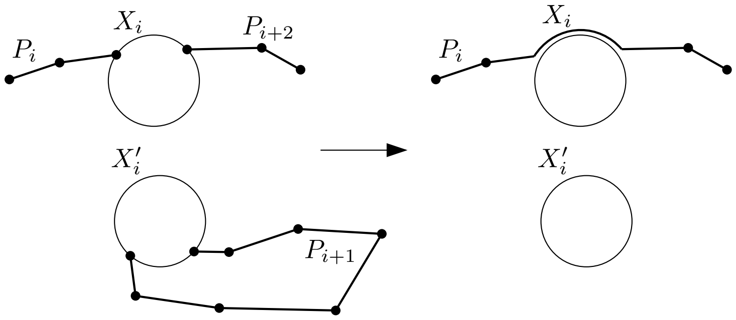

Figure 5: Reducing a path.

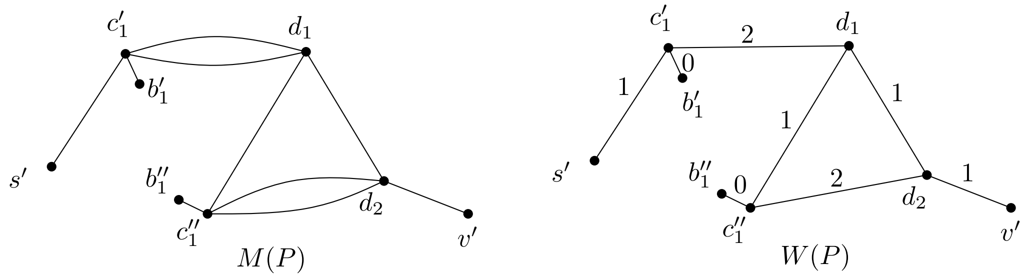

Figure 6: The characteristic multigraph of the path P from Figure 4 (left) and its corridor multigraph (right).

at each end.

We also add auxiliary leaf vertices b'_i and b''_i joined by auxiliary edges to c'_i and c''_i. The points b'_i and b''_i are chosen as points in the interiors of the cycles C'_i and C''_i, between two consecutive edges $x'_iv'_j, x'_iv'_{j+1}$ and $x''_iv''j, x''_iv''_{j+1}$, where v'_j and v''_j (resp. v'_{j+1} and v''_{j+1}) are corresponding vertices on the cycles C'_i and C''_i. Then we can define a linear ordering of the (non-auxiliary) edges incident with c'_i and c''_i, starting at b'_i (b''_i) and going around the vertex c'_i (c''_i) in clockwise direction. This allows us to reconstruct the type of the original path P from the characteristic multigraph $M(P)$, by walking along the edges of $M(P)$ in a deterministic way. The path starts at s' and ends at v'. When the path arrives at the vertex c'_i (c''_i) along the j-th leftmost edge (in the linear ordering constructed above), it continues from c''_i (c'_i) along the j-th rightmost edge. When the path arrives at d_i, it continues from the same vertex along the opposite edge incident with d_i.

If the multigraph $M(P)$ has more than one component, we sequentially add auxiliary edges joining two different components, to get a connected multigraph $M'(P)$.

Now we construct a compact description of the multigraph $M'(P)$. We define a weighted embedded undirected *corridor multigraph* $W(P)$ as follows. Put

$W_0(P) = M(P)$ and set the weights of all edges to 1 (except the auxiliary edges that get weight 0). Then repeat the following step: if there is a face f in $W_i(P)$ bounded by only two edges e_1 and e_2 with weights $w(e_1)$ and $w(e_2)$, construct $W_{i+1}(P)$ from $W_i(P)$ by merging the edges e_1 and e_2 into a single edge with weight $w(e_1) + w(e_2)$, thus removing the face f. After at most $2^g \cdot n$ steps we get a multigraph $W(P)$ where all faces are bounded by at least three edges. Indeed, $W(P)$ contains no loops as P is a reduced path, and all bigons were eliminated in the construction of $W(P)$. By Euler's formula, $W(P)$ has at most $3(2g + 2) + 6$ edges, while the sum of the weights of the edges is at most $2^g \cdot n$. Clearly, we can reconstruct $M(P)$ (and hence the type and the terminal region of P) from the corridor multigraph $W(P)$ by replacing each edge of weight w by w parallel edges.

The construction of $W(P)$ can be achieved in logspace as follows. Take the edges of $M'(P)$ in clockwise order around each vertex. If the current edge e_i forms a bigon with the previous edge e_{i-1}, increase the weight of the corresponding edge in the corridor multigraph by one. Otherwise form a new edge of weight 1 (or 0 if e_i is auxiliary) if e_i has not been visited already from its other endpoint.

Up to homeomorphism, there are only a constant number of connected plane multigraphs with at most $3(2g + 2) + 6$ edges. Thus all such multigraphs can be generated in constant space. To get all the possible corridor multigraphs, we need to assign weights to the edges. As each of the multigraphs has a constant number of edges and the sum of weights is linear in n, there are only polynomially many different assignments of the weights to the edges. Therefore, there are only polynomially many corridor multigraphs, up to homeomorphism. Clearly, all the weight assignments (decompositions of a number $m \leq 2^g \cdot n$ into a sum of constantly many nonnegative integers) can be generated in logspace.

By the discussion above, we can easily determine the type of the path corresponding to a given corridor multigraph in logspace. The graph K is now constructed by taking a copy of G' for each of the generated types, and gluing the appropriate pairs of these copies together.

Acknowledgments

We are grateful to Eric Allender and Michael Saks for introducing us to the topic, for their encouragement and many helpful discussions. We also thank Marek Sterzik for participating in the discussions in early stages of our research.

References

[1] E. Allender, D. A. M. Barrington, T. Chakraborty, S. Datta and S. Roy, Planar and grid graph reachability problems, to appear in Theory of Computing

Systems.

[2] E. Allender and M. Mahajan, The complexity of planarity testing, *Inform. and Comput.* **189**(1) (2004), 117–134.

[3] C. Bourke, R. Tewari and N. V. Vinodchandran, Directed planar reachability is in unambiguous log-space, *ACM Trans. Comput. Theory* **1**(1) (2009), 1–17.

[4] A. Jakoby, M. Liśkiewicz and R. Reischuk, Space efficient algorithms for directed series-parallel graphs, *J. Algorithms* **60**(2) (2006), 85–114.

[5] N. Nisan and A. Ta-Shma, Symmetric Logspace is closed under complement, *Chicago J. Theoret. Comput. Sci.* (1995), Article 1, 13pp. (electronic).

[6] O. Reingold, Undirected connectivity in log-space, *J. ACM* **55**(4) (2008), Art. 17, 24.

[7] T. Thierauf and F. Wagner, Reachability in $K_{3,3}$-free graphs and K_5-free graphs is in unambiguous log-space, Tech. Rep. 29, Electronic Colloquium on Computational Complexity (2009).

ECCC ISSN 1433-8092

http://eccc.hpi-web.de/