

Parallel Complexity of Depth-First-Search and Maximal Path in Restricted Graph Classes

Archit Chauhan¹ 

Department of Computer Science and Engineering, IIT Bombay, India

Samir Datta 

Chennai Mathematical Institute, India

UMI ReLaX, Indo-French joint research unit, Chennai, India

M. Praveen 

Chennai Mathematical Institute, India

UMI ReLaX, Indo-French joint research unit, Chennai, India

Abstract

Constructing a Depth First Search (DFS) tree is a fundamental graph problem, whose parallel complexity is still not settled. Reif showed parallel intractability of lex-first DFS. In contrast, randomized parallel algorithms (and more recently, deterministic quasipolynomial parallel algorithms) are known for constructing a DFS tree in general (di)graphs. However a deterministic parallel algorithm for DFS in general graphs remains an elusive goal. Working towards this, a series of works gave deterministic NC algorithms for DFS in planar graphs and digraphs. We further extend these results to more general graph classes, by providing NC algorithms for (di)graphs of bounded genus, and for undirected H -minor-free graphs where H is a fixed graph with at most one crossing. For the case of (di)graphs of bounded treewidth, we further improve the complexity to a Logspace bound. Constructing a maximal path is a simpler problem (that reduces to DFS) for which no deterministic parallel bounds are known for general graphs. For planar graphs a bound of $O(\log n)$ parallel time on a CRCW PRAM (thus in NC^2) is known. We improve this bound to Logspace.

2012 ACM Subject Classification Theory of computation \rightarrow Parallel algorithms

Keywords and phrases Parallel Complexity, Graph Algorithms, Depth First Search, Maximal Path, Planar Graphs, Minor-Free, Treewidth, Logspace

Digital Object Identifier 10.4230/LIPIcs.FSTTCS.2025.23

Related Version *Full Version*: <https://arxiv.org/pdf/2506.14974v2>

Funding *Samir Datta*: Partially supported by a grant from Infosys foundation.

M. Praveen: Partially funded by a grant from the Infosys foundation.

1 Introduction

Depth First Search (DFS) is an archetypal graph algorithm that has played a pivotal role in the study of several problems such as biconnectivity, triconnectivity, topological sorting, strong connectivity, planarity testing and embedding etc. Along with Breadth First Search (BFS) it is perhaps one of the best known algorithms. However, while the BFS algorithm is easy to parallelize, no such deterministic parallel algorithm is known for DFS². In fact,

¹ Part of this work was done when the author was a PhD student at Chennai Mathematical Institute, India.

² When talking about parallel computation, we will mean the class NC. In terms of PRAM models, a problem is said to be in the NC if there exists a PRAM machine and constants k, c such that the machine solves the problem in $O(\log^k n)$ time, using $O(n^c)$ many parallel processors. If the machine is a Concurrent Read Concurrent Write PRAM, we denote this class by $CRCW[\log^k n]$. An equivalent definition in terms of boolean circuits also exists. In particular $CRCW[\log^k n]$ is equivalent to AC^k (see [39, Section 3.4]).



Reif showed in the 1980's [48] that unless all of P is parallelizable, there is no parallel algorithm that can mimic the familiar recursive DFS. To be slightly more precise finding the lexicographically first DFS-order is P-complete. Roughly contemporaneously, the problem of finding a (not necessarily lexicographically first) DFS-order was reduced to the minimum weight perfect matching problem for bipartite graphs in [1] (and in [2] for directed graphs) which coupled with randomized parallel algorithm for min-weight perfect matching [40, 47], yielded a randomized parallel algorithm for DFS. Recent advances in matching algorithms [28] indeed yield a deterministic parallel algorithm for DFS although using quasipolynomially many (in particular, $n^{O(\log n)}$ many) processors.

A natural question to consider is: for which restricted graph classes can we give deterministic parallel algorithms for DFS. Planar graphs, as well as digraphs have been known to possess (deterministic) parallel DFS algorithms dating back to half a century ago [53, 33, 35, 38, 37, 3]. Khuller also gave an NC algorithm for DFS in $K_{3,3}$ -minor-free graphs [41]. The above mentioned algorithms differ in their parallel complexity. For example Kao, Klein in [38] gave a parallel running time of $O(\log^{10} n)$ for DFS in planar digraphs while optimizing the number of required processors to be linear. More recently, Ghaffari et al [29] have given a (randomized) nearly work efficient algorithm for DFS running in $\tilde{O}(\sqrt{n})$ time and $\tilde{O}(m + n)$ work. Our metric for evaluating an algorithm is more complexity theoretic – so we try to optimise on the complexity class we can place the problem in. In particular, we are interested in placing the problems in the smallest possible class in the following (loose) hierarchy : $L \subseteq UL \subseteq NL \subseteq AC^1 \subseteq NC^2 \subseteq AC^3 \subseteq \dots$ (See Chapter 6 of [9]).

Most algorithms for parallel DFS in graphs proceed via finding a balanced path separator in the graph, which is a path such that removing all of its vertices leaves behind connected components (or strongly connected components in case of digraphs) of small size, allowing for logarithmic depth recursion. It was shown by Kao in [35] that all graphs and digraphs have balanced path separators, and it is easy to find one sequentially using the DFS algorithm itself. The challenge thus is to find a path separator in parallel. Another useful result Kao [36, 37] also showed is that if we have a small number of disjoint paths that together constitute a balanced separator (by small we mean up to polylogarithmically many, we call these *multi-path separators*), then we can trim and merge them into a single balanced path separator in parallel. If we could somehow merge and reduce a multi-path separator consisting of a large number of paths, say $O(n)$ many, to a multi-path separator consisting of up to constant (or even polylogarithmically many) number of paths in NC, then we would immediately have an NC algorithm for DFS, as the set of all vertices of the graph constitutes a trivial multi-path separator. Such a routine to reduce the number of paths in a multi-path separator is indeed the heart of the RNC algorithm of Aggarwal and Anderson [1] (and also [2]), and this is where the randomized routine for minimum weight perfect matching in bipartite graphs is used.

1.1 Our contributions

We first look at graphs of bounded genus in Section 4, which are a natural generalization of planar graphs (planar graphs have genus zero). We are not aware of any deterministic NC algorithm for DFS even for toroidal graphs. We give an NC algorithm for DFS which proceeds via finding a multi-path separator consisting of $O(g)$ many paths, where g is the genus of the graph. For graph classes of higher genus, the main bottleneck we encounter is that of testing their genus and finding an embedding into a surface of that genus in NC. More precisely, if we assume that we can find genus and embeddings of graphs of upto $O(\log n)$ genus in NC, then our algorithm still gives an NC bound for graphs of up to $O(\log n)$ genus.

In Section 5 we look at DFS in single crossing minor free graphs, which are another generalization of planar graphs (orthogonal to bounded genus graphs). By Wagner’s theorem [58], we know that planar graphs are precisely those whose set of forbidden minors consists of K_5 and $K_{3,3}$. Both $K_5, K_{3,3}$ are single crossing graphs, i.e. they can be drawn on the plane with at most one crossing. We consider the family of H -minor-free graphs, where H is any fixed single crossing graph. Such families which include in particular the families of K_5 -minor-free graphs and $K_{3,3}$ -minor-free graphs, are called single-crossing-minor-free graphs. They form a natural stepping stone towards more general H -minor-free families, which by a deep theorem of Robertson and Seymour [51], completely characterize *minor-closed* families. The only result we know of in this direction is that of Khuller’s [41] parallel algorithm for DFS in $K_{3,3}$ -minor-free graphs. We generalize this result by giving an NC algorithm for DFS in single-crossing-minor-free graphs. The algorithm of [41] uses a decomposition of $K_{3,3}$ -minor-free graphs by 2-*clique sums* (or separating pairs) into pieces that are either planar or are exactly the graph K_5 . They proceed by finding the “center” piece of this decomposition tree and then finding a balanced cycle separator in it, after projecting the weights of attached subgraphs on the center piece. The high level strategy of our algorithm is similar. We first use a decomposition theorem [50] by Robertson and Seymour which states that single-crossing-minor-free graphs can be decomposed by 3-*clique sums* into pieces that are either planar or of bounded treewidth. Then we try to find a balanced cycle separator in the center piece. However the attachment at 3-cliques instead of 2-cliques poses some technical difficulties, as simply projecting the weights of attached subgraphs on the center piece does not work. To deal with this, we use appropriate gadgets while projecting, that mimic the connectivity of the attached subgraphs. We remark here that the problem of finding a minimum weight perfect matching was shown to be in (deterministic) NC for bounded genus bipartite graphs [19, 32] as well as for single-crossing-minor-free graphs [27]. However, the reductions in [1, 2] involve adding large bi-cliques in the input graph, which does not seem to preserve the genus, nor the forbidden minors of the graph.

In Section 6 we look at bounded treewidth (di)graphs, for which we are not aware of any explicitly stated parallel algorithms for DFS. We note however that the machinery of Kao [35] easily yields an $AC^1(L)$ (which is contained in AC^2 . See Section 2.) algorithm for DFS in bounded treewidth graphs. We improve this bound to L . Instead of using path separators, our technique for bounded treewidth (di)graphs involves the parallel (in fact, Logspace) version of Courcelle’s theorem from [25]. We also use an existing result [17, 11] that states that we can add a linear ordering to the vertices of the graph without blowing up the treewidth of the structure. This combined with an observation of [23] characterizing the lex-first DFS tree in terms of lex-first paths to each vertex, allows us to express membership of an edge to the DFS tree, in MSO.

Finally, in Section 7 we turn to the problem of finding a maximal path in a planar graph, or **MaximalPath**. Given an input (undirected) graph and a vertex r in it, the problem is to find a (simple) path starting at r that cannot be extended further, i.e., all the neighbours of its end point other than r already lie on the path. This problem easily reduces to DFS because every root to leaf path in a DFS-tree is a maximal path. However even for this problem, no deterministic parallel algorithm for general graphs is known, and the RNC algorithm by [7] goes via reduction to perfect matching. For restricted graph classes, the situation is clearer for **MaximalPath** as compared to DFS. Anderson and Mayr [8] gave a deterministic $O(\log^3 n)$ time algorithm on a CREW PRAM, for graphs of bounded degeneracy³ These include for

³ A graph is said to be k -degenerate if every subgraph of it has a vertex of degree at most k .

example, planar graphs, and even more generally, all H -minor free graph classes. In modern terms, their algorithm witnesses a bound of $\text{AC}^1(\text{L})$. A better bound of $\text{CRCW}[\log n]$ for **MaximalPath** in *planar graphs* comes from the DFS algorithm of Hagerup [33], which again can be improved by a modern analysis to $\text{UL} \cap \text{co-UL}$ [4]. We improve this bound to L . Our algorithm proceeds by first reducing the problem to finding a maximal path in the triconnected components of a planar graph using the 2-clique sum decomposition (also referred to as the **SPQR** decomposition) of [21]. We then show that in the triconnected planar components, a maximal path can be found in the “outerplanar layers” of the graph. The results we prove can be summarized as follows:

► **Theorem 1.** *We have the following bounds:*

1. DFS in bounded genus graphs and digraphs is in $\text{AC}^1(\text{UL} \cap \text{co-UL})$.
2. DFS in undirected single-crossing-minor-free graphs is in NC .
3. DFS in directed and undirected graphs of bounded treewidth is in L .
4. **MaximalPath** in undirected planar graphs is in L .

2 Preliminaries

We assume the reader is familiar with tree-decompositions, planarity, embeddings on surfaces and graph minors. When referring to concepts like treewidth, minors and genus of digraphs, we will always intend them to apply to the underlying undirected graph. In digraphs, we will use the term path to mean simple directed path unless otherwise stated. We will use *tail* of a path to refer to the starting vertex of the path and *head* of a path to refer to the ending vertex of the path. The term *disjoint paths* will always mean vertex disjoint paths, and the term *internally disjoint paths* will refer to paths that are disjoint except possibly sharing end points. In a given graph $G = (V, E)$ and edge $(u, v) \in E(G)$, we denote the graph obtained by removing the edge (u, v) (but keeping the vertices u, v intact) by $G - (u, v)$.

We refer the reader to [9] for definitions of classes $\text{L}, \text{NL}, \text{AC}, \text{NC}$. A nondeterministic Turing machine is said to be unambiguous if for every input, there is at most one accepting computation path. The class UL is the class of problems for which there exists an unambiguous Logspace Turing machine to decide the problem. In our setting, our outputs are not binary, but encodings of objects like a DFS tree, or a maximal path. Each of the classes mentioned above can be generalised to deal with *functions* instead of just *languages* in a natural way. One may refer to [34] or [3, Section 2] for more details, specially what is meant by the computation of functions by a bounded space class like L, NL or $\text{UL} \cap \text{co-UL}$. In simplest terms, the output has a size of polynomially many bits and each bit can be computed in the stated class. We will often refer to a machine that computes a function as a transducer. A constant number of L transducers may be composed in L (see [9, Chapter 4]). A similar result holds for classes like $\text{NL}, \text{UL} \cap \text{co-UL}$ (see [3, Section 2] for a more elaborate explanation).

An important result in the Logspace world is that of Reingold [49] which states that the problem of deciding reachability (as well as its search version) in undirected graphs is in L . Using this and the lemma of composition of L transducers, many routines like computing connected components left after removing some vertices, or counting the number of vertices in a connected component, can be shown to be in L . Another result we will use is that the problem of reachability in directed planar graphs is in $\text{UL} \cap \text{co-UL}$, as shown by [13, Theorem 1.1]. This extends to graphs of bounded genus, and in fact to graphs of $O(\log n)$ genus if their polygonal schema is given as part of input [31, Theorem 7] (See also [43, 54, 19]). The theorem of [31] achieves this by giving a logspace algorithm to compute $O(\log n)$ size weight functions for edges that isolate shortest paths between vertices of the input graph. In

other words, for any pair of vertices u, v , there is a unique path minimum weight path from u to v under the weight function computed by [31, Theorem 7]. Thierauf and Wagner [55] showed that with such path isolating weights assigned to edges, we can compute distances as well as shortest path between vertices, in $\text{UL} \cap \text{co-UL}$. These results will be useful in computing arborescences⁴ in bounded genus graphs in Section 4.

Suppose \mathcal{C} is a complexity class (a set of boolean functions). We use $\text{AC}^k(\mathcal{C})$ to denote functions that can be computed by a family of AC^k circuits augmented with oracle gates that can compute functions in \mathcal{C} . Note for example, that since $\text{UL} \cap \text{co-UL} \subseteq \text{AC}^1$, the class $\text{AC}^1(\text{UL} \cap \text{co-UL})$ is contained in AC^2 .

We will use the notion of Gaifman graph of a logical structure:

► **Definition 2.** *Given a logical structure \mathcal{A} , with a finite universe $|\mathcal{A}|$, and relation symbols which are interpreted as relations of \mathcal{A} , the Gaifman graph of \mathcal{A} , $G_{\mathcal{A}}$ is a graph whose vertices are the elements of $|\mathcal{A}|$, and for $u, v \in |\mathcal{A}|$, we add an edge between them iff u, v both occur in a relation tuple of \mathcal{A} .*

The treewidth of the structure is the treewidth of its Gaifman graph. We will use the following extension of Courcelle's theorem by [25]:

► **Theorem 3** ([25]). *For every $k \geq 1$ and every MSO-formula ϕ , there is a logspace DTM that on input of any logical structure \mathcal{A} of treewidth at most k , decides whether $\mathcal{A} \models \phi$ holds.*


Clique-Sum decompositions

We recall the notion of *clique-sums*. Given two graphs G_1, G_2 , a k -clique-sum of them can be obtained from the disjoint union of G_1, G_2 by identifying a clique in G_1 of at most k vertices with a clique of the same number of vertices in G_2 , and then possibly deleting some of the edges of the merged clique. It is used in reverse while decomposing a graph by clique sums. For our purpose, we will use up to 3-clique sum decomposition of a graph. Suppose G decomposes via a 3-clique sum at clique c into G_1 and G_2 . Then we write G as $G_1 \oplus_c G_2$. More generally, if G_1, G_2, \dots, G_ℓ all share a common clique c , then we use $G_1 \oplus_c G_2 \oplus_c \dots \oplus_c G_\ell$ to mean G_1, G_2, \dots, G_ℓ are glued together at the shared clique. When separating the graph along a separating pair/triplet, we add virtual edges if needed to make the separating pair/triplet a clique. The decomposition can be thought of as a two colored tree (see [14, 27, 16]) where the blue colored nodes represent *pieces* (subgraphs that the graph is decomposed into), and the red colored nodes represent the cliques at which two or more pieces may be attached. The edges of the clique describe the incidence relation between pieces and cliques (see Figure 2). For single-crossing-minor-free graphs, Robertson and Seymour gave the following theorem regarding their 3-clique sum decomposition:


► **Theorem 4** ([50]). *For any single-crossing graph H , there is an integer τ_H such that every graph with no minor isomorphic to H can be obtained by 3-clique sums, starting from planar graphs and graphs of treewidth at most τ_H .*

An observation we use is that we can assume that in any planar piece of the decomposition, the vertices of a separating pair or triplet lie on a common face (Otherwise we could decompose the graph further. See [14, 27] for a more detailed explanation). It is shown in [27] that the decomposition can be computed in NC. For computing the 2-clique sum decomposition, we

⁴ An arborescence of a digraph G is a spanning tree rooted at a vertex r , such that all edges of the tree are directed away from r .



■ **Figure 1** An example of a graph G . We ignore directions here.



■ **Figure 2** A clique sum decomposition of G . Red nodes are the clique nodes and blue node the piece nodes. Dashed edges denote virtual edges.

will use the Logspace algorithm given by [21]. The resulting pieces of the decomposition, also called as the *triconnected components* of the graph, are either dipoles, simple cycles, or non-trivial 3-connected graphs (i.e. 3-connected graphs with at least four vertices).

Balanced Path/Cycle separators

Consider an input graph $G = (V, E)$ with $|V| = n$. If G is undirected, we use the term α -separator (or just balanced separator) in general to mean some subgraph of G which if removed from G , then the size of each of the resulting connected components is at most αn , where $0 < \alpha < 1$. A single path P , or a collection of paths $\{P_1, \dots, P_k\}$ for example could form a balanced separator of G . We call the latter a *multi-path* separator consisting of k paths. If G is directed, we have a similar definition, but the constraint is instead on the size of each *strongly* connected component left after removing the vertices of the separator. Note that if we have an algorithm to find an α -separator in a directed graph, then we can use that to find an α -separator in an undirected graph by making each edge into a bi-directed edge. Any simple cycle C in an embedded planar graph G will divide it into two regions, the *interior* and *exterior* of C , which we denote by $int(C)$ and $ext(C)$ respectively. Typically, if a face of G is marked as the outer face then the region containing the outer face is referred to as the exterior and the other region as the interior. Suppose the vertices, edges and faces of G are assigned weights. Then we use $w(int(C))$ to denote the sum of weights of vertices, edges, faces that are contained in $int(C)$. The term $w(ext(C))$ is defined similarly. The vertices, edges that lie on C itself are usually not counted since if they are part of the separator and will be removed when recursing. Suppose u, v are two vertices in the cycle C (embedded in the plane). We denote by $u \circlearrowright v$ (respectively $u \circlearrowleft v$), the segment of C obtained by starting from u and going clockwise (respectively counter-clockwise) along C till v .

3 Path/Cycle Separators review

We restate some of the existing algorithms and analyze their complexity with respect to circuits/space bounded complexity classes. First, we reiterate an algorithm of Kao (see also [35, 2, 37]) that combines a small number of path separators into a single path separator in parallel, and observe that it is in L given an oracle for REACH. We state the theorems and corollaries for digraphs, but the corresponding statements for undirected graphs also hold with appropriate changes.

► **Theorem 5** ([36, 37]). *Let G be a digraph and $\{P_1, P_2\}$ be two paths that together form an α -separator of G for some $\alpha \in [\frac{1}{2}, 1)$. Given an oracle for REACH, we can in L , construct a single path P_3 that is also an α -separator of G .*

We reproduce the proof for completeness in Section A.1.1. Note that as a corollary, the above algorithm is in NL for general directed graphs, and in $UL \cap co-UL$ for directed graphs of $O(\log n)$ genus (using [31]). For undirected graphs, the algorithm (with appropriate changes) is in L using Reingold's algorithm [49]. As noted in [2, 38], if we have a multi-path separator consisting of k many paths, say $\{P_1, P_2 \dots P_k\}$, we can combine them into a single path separator using k iterations of the above procedure with k transducers. Thus we have:

► **Corollary 6.** *Let G be a digraph and k be a constant. Let $\{P_1, P_2, \dots P_k\}$ be k disjoint paths that form a multi-path α -separator of G for some $\alpha \in [\frac{1}{2}, 1)$. Given an oracle for REACH, we can in L , construct a single path P_{k+1} that is also an α -separator of G .*

For bounded treewidth graphs, we can find a tree decomposition in L using [25, Theorem 1.1], and we know that there is always a bag in the tree decomposition of a graph whose vertices together form a $1/2$ -separator of the graph. Since REACH is known to be in L for digraphs of bounded treewidth (from Theorem 3), we have the following result:

► **Corollary 7.** *Given a digraph of bounded treewidth, we can find a $1/2$ -path separator in L .*

It is shown in [3, Section 7] how to construct a DFS tree in a digraph given a routine for finding path separators, in parallel. The algorithm follows a divide and conquer strategy, where we find DFS trees on the smaller strongly connected components in parallel, and sew them together using a routine for DFS in DAGs. We state the result in the following proposition.

► **Proposition 8.** *Let G be a digraph and $r \in V(G)$. Let $F_{sep}(H)$ be a function that takes as input any digraph H and returns a balanced path separator of H . Then there exists a uniform family of AC^1 circuits with oracle gates for F_{sep} and functions computable in $UL \cap co-UL$, that take as input (G, r) , and return a DFS tree of G rooted at r .*

Although the proposition is for digraphs, the same idea works in the simpler setting of undirected graphs also (Instead of DAGs of strongly connected components, we just have to recurse on connected components there. See for example [53]). We will also use the following, more general notion of cycle separators in planar graphs, that was introduced in [45]:

► **Definition 9.** *Let G be an undirected planar graph, and $w: \{V(G) \cup E(G) \cup F(G)\} \rightarrow \mathbb{Q}^+$ be a normalised weight function assigning weights to vertices, edges and faces of G that sum up to 1. We call a cycle C a $2/3$ - interior-exterior cycle separator if $w(int(C)) \leq 2/3$, and $w(ext(C)) \leq 2/3$.*

Miller showed in [45] that such a weighted cycle separator always exists if G is biconnected and if no face is too heavy. Moreover, we can find such a separator in parallel⁵ (see also [52]).

► **Theorem 10** ([45]). *Let G be a biconnected planar graph, and $w: \{V(G) \cup E(G) \cup F(G)\} \rightarrow \mathbb{Q}^+$ be an assignment of non-negative weights to vertices, edges, faces that sum up to 1, such that no face has weight more than $2/3$. Then we can construct a $2/3$ -interior-exterior-cycle separator of G in $\text{CRCW}[\log n]$.*

4 DFS in directed graphs of bounded genus

We show how to find balanced path separators (and therefore a DFS tree) in digraphs that can be embedded on a surface of genus at most a fixed constant g , in NC. For an undirected graph, it is easy to see that a DFS tree of the bidirected version of it is also a DFS tree of the original graph.

The first step naturally is to find an embedding, which can be done using the theorem of Elberfeld and Kawarabayashi [26, Theorem 1], which states that we can find an embedding of a graph of genus at most a constant g , on a surface of genus g , in Logspace. The idea is to find a small number of disjoint, surface non-separating cycles (not necessarily directed cycles), say C_1, C_2, \dots, C_l ($l \leq g$) in G , such that the (weakly) connected components left after removing them are all planar. These can be found with at most g sequential iterations. Then, if any of the remaining planar components has a large strongly connected component, we can find a $2/3$ -path separator, say P , in it using [38, Theorem 19]. A more modern analysis of the algorithm in [38] gives a $\text{UL} \cap \text{co-UL}$ bound (see [15, Section 3.3] for a proof. See also [3, Lemma 50], which gives an $11/12$ -path separator with the same bound). Then C_1, C_2, \dots, C_l, P together form a balanced multi-path separator of the digraph G . The (possibly undirected) cycles C_1, C_2, \dots, C_l will have an additional property. Any $C_i \in \{C_1, C_2, \dots, C_l\}$ will either be a directed cycle, or it will consist of two internally disjoint (directed) paths that share a common tail vertex and a common head vertex. Therefore the multi-path separator can be written as consisting of at most $2g + 1$ many paths, and we can use Corollary 6 to get the required path separator. We use the following known fact (see [46, Lemma 4.2.4] or [30, Fact 9.1.6]).

► **Theorem 11** ([46]). *Let G be a graph embeddable on a surface S , of genus g . Let C be a cycle that forms a surface non-separating curve on some embedding of G on S . Then every connected component of $G - C$ is embeddable on a surface of genus $g - 1$.*

Therefore, if we have an algorithm to find such a (undirected) cycle in NC, then we can repeat it g times to find the required (undirected) cycles C_1, C_2, \dots, C_k . We will use the following lemma from [5] (See also [56]) to find such cycles.

► **Lemma 12** ([5]). *Let G be a graph of genus $g > 0$, and let T be a spanning tree of G . Then there is an edge $e \in E(G)$ such that $T \cup e$ contains a surface non-separating cycle.*

Clearly the theorem also works if we have an arborescence instead of a spanning tree. The fundamental cycle C corresponding to a non-tree edge $e = (u, v)$ would consist of the non-tree edge e and two vertex disjoint directed paths, say P, P' , starting from a common ancestor of (u, v) in the arborescence (we can include the common ancestor vertex in either of P or P'), and ending at u, v respectively. We can find an arborescence in $\text{UL} \cap \text{co-UL}$ using [31, 55] (as

⁵ The theorem in [45] also gives a bound on size of the separator but that is not important for us.

explained in the appendix), and the required fundamental cycle in Logspace using the theorem of [26]. We describe the steps of the algorithm in detail in Section A.2. Therefore we can find a $2/3$ -path separator in directed graphs of bounded genus in $\text{UL} \cap \text{co-UL}$. Using Proposition 8, we can find a depth first search tree in such graphs in $\text{AC}^1(\text{UL} \cap \text{co-UL})$, finishing the proof of part 1 of Theorem 1.

5 DFS in Single-Crossing-Minor-Free Graphs

We start with computing the 3-clique-sum decomposition by the NC algorithm of [27]. We will denote this decomposition tree by T_G , each node of which is either a piece node or a clique node, and the weight of each node is the number of vertices in the corresponding piece/clique. (Note that by this convention, the sum of weights of all nodes of T_G will sum up to more than n since some vertices would occur in multiple pieces).

► **Definition 13.** Let t_R be a node of T_G , and R its corresponding piece/clique, such that the size of any connected component in $G - R$ is less than $n/2$. We call t_R a central node of T_G , and R a central piece/clique.

Such a node always exists and can be found in Logspace by traversing along the heaviest child in T_G (see Chapter 7 of [18] for example). We will assume t_R to be the root of T_G hereafter. We will denote the subtrees of T_G that are children of t_R by $\{T_1, T_2, \dots, T_l\}$, the subgraphs of G corresponding to these subtrees by $\{G_1, G_2, \dots, G_l\}$, and the cliques by which they are attached to R by $\{c_1, c_2, \dots, c_l\}$ respectively. These subgraphs might themselves consist of smaller subgraphs glued at the common clique. For example, G_1 might consist of $G_{11}, G_{12}, \dots, G_{1k}$ that are glued at the shared clique c_1 (i.e. $G_1 = G_{11} \oplus_{c_1} G_{12} \dots \oplus_{c_1} G_{1k}$). See Figure 3 for reference. Note that because R is a central node, $w(G_{1i}) - w(c_1) < n/2 \forall i \in [1..k]$. If t_R is a clique node, then we have a $1/2$ -separator of at most three vertices and we are done. Therefore there are two cases to consider. Either R is a planar piece, or a piece of bounded treewidth.

When R is of bounded treewidth

In this case, we will refine the tree T_G by further decomposing R . Since R is of treewidth at most τ_H , we can compute a tree decomposition of R such that every bag is of size at most $\tau_H + 1$, in L using [25]. Let this tree be denoted by T_R . Consider the subtree T_1 of T_G as described above, which has a node t_{c_1} by which it is attached to R . Since c_1 is a clique, there must be at least one bag in T_R that contains all vertices of c_1 (see [12]). Attach t_{c_1} to any such node of T_R . Do this procedure for all subtrees T_1, T_2, \dots, T_l . It is easy to see that this will result in another tree decomposition, and that at least one of the nodes of T_R will be a central node of this new tree. Hence we can use the procedure described above to find it in L. Since the bags of T_R are of size at most $\tau_H + 1$, we get a $1/2$ -separator of G consisting of constant number of vertices, and we are done.

When R is planar

We start with a simple idea, but encounter an obstacle. Then we show how to fix it. Note that R (along with its virtual edges), is biconnected. To initialize, we assign weight 1 to all vertices of R and weight 0 to all edges and faces of R . We project the weight of subgraphs G_1, G_2, \dots, G_l on vertices, edges and faces of R , at the respective cliques they are attached at. Suppose the subgraph G_k is attached to R at clique c_k . Then we project the weight as follows.

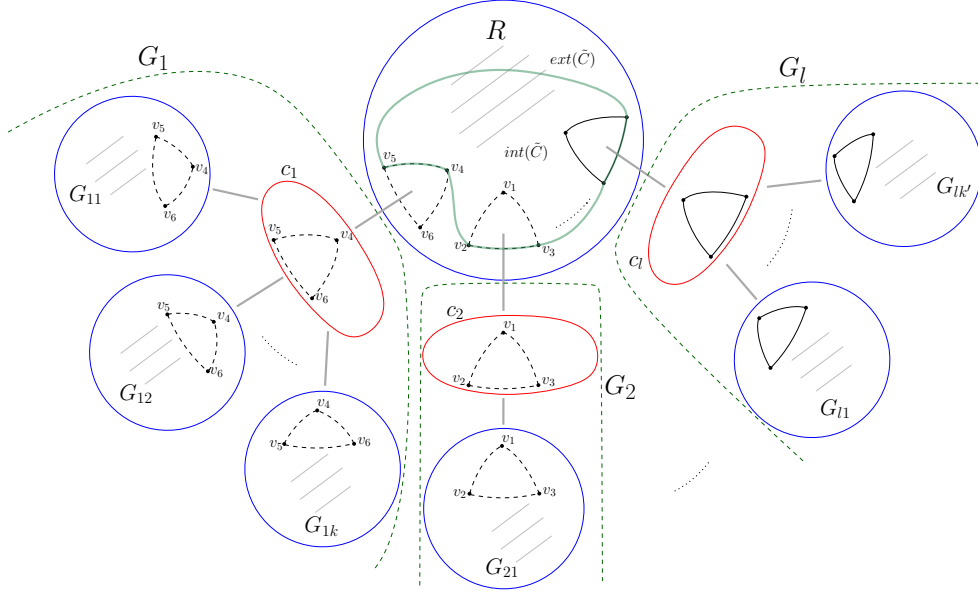


Figure 3 R is the central node with children subgraphs G_1, G_2, \dots, G_l attached to it via cliques c_1, c_2, \dots, c_l respectively. G_1 consists of subgraphs $G_{11}, G_{12} \dots G_{1k}$ glued at c_1 . G_2 is same as G_{21} . Dashed edges denote virtual edges in R . In the case when R is a planar piece, it is useful to think of it as a sphere, with separating triplets forming empty triangular faces in R , on which G_1, G_2, \dots, G_l are attached. The cycle in green is an example of *interior-exterior-cycle separator* \tilde{C} (though gadget replacements are not shown in the figure). Of the attached subgraphs shown, G_2, G_l are *interior components* with respect to \tilde{C} whereas G_1 is an *exterior component* with respect to \tilde{C} .

1. If c_k is a single vertex v , then assign a weight equal to $|V(G_k)|$ to *vertex* v .
2. If c_k is a separating pair (u, v) , then we assign a weight equal to $|V(G_k)| - 2$ to the *edge* (u, v) of R . (We subtract 2 since weights of vertices u, v are already accounted for).
3. If c_k is 3-clique of vertices (v_1, v_2, v_3) , then we assign weight equal to $|V(G_k)| - 3$ to the *face* of R enclosed by c_k .

Let \tilde{R} denote the weighted version of R obtained after projecting these weights. The sum of weights of all vertices, edges, faces of R is n . There are two possibilities. Suppose there exists a clique, say c_1 (which is either a vertex, or an edge, or a face of \tilde{R}), with weight at least $2n/3$. This means that total weight of subgraphs attached to R via c_1 , i.e. $w(G_{11} \oplus_{c_1} G_{12} \oplus_{c_1} \dots \oplus G_{1k})$, is at least $2n/3$. Since each of the graphs $G_{11}, G_{12} \dots G_{1k}$ is of size at most $n/2$ and they all get disconnected on removal of c_1 , the vertices of c_1 form a $1/2$ -separator of G , and we are done. Therefore the remaining possibility we need to consider is one where every vertex, edge, face of \tilde{R} has weight smaller than $2n/3$. Since \tilde{R} is planar and biconnected, we can find a $2/3$ -interior-exterior-cycle-separator of it in L using Theorem 10. Let \tilde{C} be such a separator. Label one of the regions \tilde{C} divides R into as the interior of \tilde{C} and the other as the exterior of \tilde{C} . Observe that for any subgraph that is attached to R (like G_2 for example), its vertices of attachment to R (the clique via which it is attached), will either all lie on or in the interior of \tilde{C} , or all lie on or in the exterior of \tilde{C} . We call the attached subgraphs of the former kind *interior components* of R with respect to \tilde{C} , and the latter as *exterior components* of R with respect to \tilde{C} (see Figure 3). We will drop the phrase “with respect to \tilde{C} ” for brevity. On removing the vertices of \tilde{C} in G , there cannot remain any path in G from any vertex of an interior component of R to any vertex of an exterior component of R . Therefore we have the following claim:

▷ **Claim 14.** Let \tilde{C} be an α -interior-exterior-cycle-separator of \tilde{R} , where $\alpha \in [\frac{1}{2}, 1)$. Then the vertices of \tilde{C} form an α -separator of G .

Now if we could replace the virtual edges of \tilde{C} by valid path segments in the attached subgraphs then we would get a simple cycle C in G which would be a $2/3$ -separator of G , and we would be done. However there is an issue here, as the virtual edges in a 3-clique need not capture “disjointness” of the external connections accurately. For example, suppose $c_2 = \{v_1, v_2, v_3\}$ is a 3-clique and the subgraph G_2 is attached to R at c_2 (as shown in Figure 3). Suppose (v_1, v_2) , and (v_1, v_3) are virtual edges in R , and our cycle separator \tilde{C} uses both of these edges (they must necessarily be consecutive in \tilde{C}). Though there must be paths between v_1, v_2 , and between v_1, v_3 in G_2 (since it is attached via 3-clique), there might not exist such paths that are internally disjoint. Hence instead of just projecting the weight of G_2 on the face formed by c_2 in R , we use a more appropriate gadget to mimic the connectivity of G_2 .

Gadgets for the subgraphs attached to R via 3-cliques

The gadget for an attached graph like G_2 is essentially a coarser planar projection of its block-cut tree, that captures the needed information regarding the cut vertices associated with its terminals $\{v_1, v_2, v_3\}$. We can adversarially assume that all three edges (v_1, v_2) , (v_2, v_3) , (v_3, v_1) are virtual edges.

► **Definition 15.** Let G_2 be a graph and vertices $v_1, v_2, v_3 \in V(G_2)$ be called its terminals. The disjoint path configuration of G_2 with respect to its terminals v_1, v_2, v_3 consists of three boolean variables : $DP(v_1-v_2-v_3)$, $DP(v_2-v_3-v_1)$, $DP(v_3-v_1-v_2)$. $DP(v_1-v_2-v_3)$ takes value True if there exists a path between v_1 and v_3 that goes via v_2 , and False otherwise. $DP(v_2-v_3-v_1)$, $DP(v_3-v_1-v_2)$ are defined similarly.

There are four cases we consider for possible disjoint path configurations, others are handled by symmetry.

Case 1 :

$DP(v_1-v_2-v_3)$	$DP(v_2-v_3-v_1)$	$DP(v_3-v_1-v_2)$
True	True	True

This means G_2 is biconnected, and hence we do not need a new gadget. We just keep the virtual edges between the clique vertices and assign weight equal to $|V(G_2)| - 3$ to the face enclosed by them (as described in Item 3). If \tilde{C} takes two virtual edges of c_2 , we can use the two disjoint path algorithm in [42] to find the corresponding paths in G_2 .

Case 2 :

$DP(v_1-v_2-v_3)$	$DP(v_2-v_3-v_1)$	$DP(v_3-v_1-v_2)$
True	True	False

By Menger’s theorem (see [44], or Chapter 3 of [24]), we know that there must be at least one cut vertex in G_2 that separates $\{v_1\}$ from $\{v_2, v_3\}$. Let x be a cut vertex separating these. Let the connected component of $G_2 - x$ containing v_1 , augmented with x , be H_1 . Let the component containing v_2, v_3 , augmented with x , be H_0 (i.e. $V(H_1) \cap V(H_0) = \{x\}$ and $V(H_1) \cup V(H_0) = V(G_2)$). We choose x such that there is no cut vertex in H_0 , that separates $\{v_1\}$ and $\{v_2, v_3\}$ in G_2 (i.e. x is the cut vertex described that is “furthest” from v_1). It is easy to see that x is unique. By our choice of x , there must be a path in H_0 from v_2 to v_3 via x . Thus the gadget G'_2 is as shown in Figure 4, containing vertices v_1, v_2, v_3, x , and the total weight distributed among the faces as shown. The following lemma shows the correctness of this procedure.




Figure 4 Figure (a) shows the gadget for graph G_2 for Case 2, (b) for Case 3 and (c) for Case 4. The faces formed by cycles of the gadget corresponding to the subgraphs H_0, H_1, H_2, H_3 of G_2 are assigned the weight functions as shown. The total weight in each gadgets sums up to $w(G_2)$. The subgraph H_0 of G_2 has no cut vertex in each case. In the gadget for Case 4, it is possible that the block H_0 is empty, i.e. there exists a vertex ($x = y = z$) in G_2 which on removal disconnects all of v_1, v_2, v_3 from each other. We don't draw that gadget separately since its structure is clear.

Lemma 16. Suppose G_2 , which is the subgraph attached to R via c_2 has the disjoint path configuration as described in Case 2. Suppose it is replaced in \tilde{R} by the gadget G'_2 defined above for this case. Then if an α -interior-exterior-cycle separator \tilde{C} of \tilde{R} uses any edges of this gadget, we can replace them by paths in G_2 such that the new cycle C is a simple cycle and it remains an α -separator in $R \oplus G_2$.

Proof. We defer the proof to the full version of this paper.

Case 3 :

$DP(v_1-v_2-v_3)$	$DP(v_2-v_3-v_1)$	$DP(v_3-v_1-v_2)$
False	True	False

The gadget is shown in Figure 4. As in the previous case, we first find the “furthest” vertex x separating $\{v_2, v_3\}$ and $\{v_1\}$. Then, in the component containing v_2, v_3 (augmented with x), we find the cut vertex y separating $\{x, v_3\}$ and $\{v_2\}$, that is furthest from v_2 .

Case 4 :

$DP(v_1-v_2-v_3)$	$DP(v_2-v_3-v_1)$	$DP(v_3-v_1-v_2)$
False	False	False

Like in previous cases, we find the “furthest” cut vertices x, y, z that separate v_1, v_2, v_3 from rest of graph. There are two cases, either there exists a single cut vertex w such that $G_2 - w$ disconnects all of v_1, v_2, v_3 from each other, which is essentially the case when $x = y = z = w$. Other case is when there is no such cut vertex. The gadget is shown in Figure 4. The proof of correctness of gadgets for Cases 3,4 is similar to the proof of Lemma 16 presented in the full version.

To find these cut vertices, we can in parallel check if each vertex of G_2 satisfies the properties of vertices “ x, y, z ” described above. Thus the gadgets can be constructed in NC. We described the procedure for G_2 , but the construction can be done in parallel for all subgraphs attached to R via 3-cliques. After finding a $2/3$ -interior-exterior-cycle separator of \tilde{R} separator using Theorem 10, we can, in parallel, replace the virtual/gadget edges by actual paths in the attached subgraphs to get the cycle separator of G . In the blocks of gadgets where we need to find a path between two vertices via a third vertex, we can use the 2-disjoint path algorithm described in [42]. Thus we can construct a $2/3$ -path-separator of G in NC. Combined with Proposition 8, this finished the proof of part 2 of Theorem 1.

6 DFS in bounded treewidth directed graphs in L

The input graph G can be seen as a structure \mathcal{A} whose universe, $|\mathcal{A}|$ is $V(G)$, and its vocabulary consists of binary relation symbol E which is interpreted as the edge relation $E(G)$ over $V(G)$. The Gaifman graph of \mathcal{A} , $G_{\mathcal{A}}$, is isomorphic to G . We will need to add a linear order O on $|\mathcal{A}|$, which would amount to adding $|V(G)|$ many edges in $G_{\mathcal{A}}$. A known result in model theory (see [11, Lemma 1], also [17, Theorem VI.4]) is that we can add a linear order to a structure, such that the treewidth of the structure is increased by at most a constant. It is shown in [10, Lemma 5.2] that we can construct such a linear order in L. Therefore after computing the order O and adding it to \mathcal{A} the new structure \mathcal{A}' also has bounded treewidth. Moreover, for any two vertices u, v , we can query if u is “lesser than” v in the transitive closure of O , in logspace. We also add the edges $E(G)$ to our universe of discourse, as well as an incidence relations : $tail(e, v)$ which is true iff the edge $e \in E(G)$ has vertex v as its tail, and $head(e, v)$ defined similarly. This can again be done without blowing up the treewidth (see section 6 of the full version).

We use $u <_O v$ to denote that vertex u is lesser than vertex v in the transitive closure of O . An ordering over vertices defines a unique lexicographically first DFS tree. The algorithm will go over each edge of E , and use Theorem 3 to query if it belongs to the lex-first DFS tree of G , in L. Our universe is $V \cup E$, and the relations on it given are $tail$, $head$, and the linear order O . To express the membership of an edge in the lex-first DFS tree with respect to O , we will use a theorem connecting the lex-first DFS tree to lex-first paths. We define lexicographic ordering on paths starting from a common vertex:

► **Definition 17.** Let P_1, P_2 be two paths in G starting from r . We say that P_1 is lexicographically lesser than P_2 (with respect to O) if at the first point of divergence starting from r , P_1 diverges to a vertex that is lesser in the given ordering than the one P_2 diverges to. We denote this as $P_1 <_O P_2$.

This naturally defines the notion of the lexicographically least path starting from r and ending at a vertex v . We call it the lex-first path from r to v . We use the following result which is shown in the proof of Theorem 11 in [23]:

► **Theorem 18.** Let T_l be the lex-first DFS tree of G with respect to a given linear order O . Then $\forall v \in V(G)$, the unique path from r to v in T_l is exactly the lex-first path from r to v (with respect to O) in G .

Thus in order to check if an edge (u, v) belongs to T_l (rooted at r), it suffices to check if it is the last edge of the lex-min path in G from r to v . We define a formula $\phi_{DFS}(u, v)$, which is true iff the edge (u, v) is part of the lex-first DFS tree T_l . Thus the logspace algorithm is the following : For each edge (u, v) in E , query $\phi_{DFS}(u, v)$ and add to output tape iff the query returns yes. A DFS numbering or order of traversal of the vertices of T_l can easily be obtained from T_l by doing an Euler tree traversal of T_l with respect to the ordering O . Now all that remains is to express the formula $\phi_{DFS}(u, v)$ in MSO. We defer that to the full version.

7 Maximal Paths in planar graphs in L

In this section we will show that given a planar undirected graph G and a vertex r in it, we can find a maximal path in G starting from r , in L. We will assume that our graph is embedded in a plane with a face marked as outer face (a planar embedding can be found in

L by [6, 22]). The first step is to reduce the problem to that of finding a maximal path in a triconnected component of G that is 3-connected (i.e. it is not a cycle or a dipole). We decompose the graph into its triconnected components, which can be done in L as described in [20]. We can choose a leaf piece L of this decomposition tree which is connected to its parent piece H by separating pair, say (r_0, r_1) . The basic idea is to find a path from r (L is chosen such that it does not contain r) to one of the vertices of the separating pair, say r_0 , without going through r_1 . Then it suffices to find a maximal path in L starting from r_0 , that does not end at r_1 . Some care must be taken in the reduction however, and cases like when L is a simple cycle need to be handled. We defer the details of the reduction to the full version.

7.1 Maximal Paths in 3-connected components

By the reduction above, we can assume that our input graph is a 3-connected component of G , say L , and our goal is to find a maximal path in L starting at r_0 and not ending at r_1 . Note that r_0, r_1 lie on a common face of L , and the edge (r_0, r_1) might be a virtual edge in L . Hence we will in general look at the graph $L' = L - (r_0, r_1)$, that is, the graph obtained from L by removing the virtual edge (r_0, r_1) (not the vertices). Let the outer face of L , on which both r_0, r_1 lie, be called f_0 . The outer face of L' is naturally obtained by merging f_0 with the other face adjacent to (r_0, r_1) . Let this be called f'_0 . We first note the following observation:

► **Observation 19.** *The boundary of the face f'_0 of L' is a simple cycle.*

If not, there would be a cut vertex x in L' . Then (r_0, x) would form a separating pair of L , contradicting the fact that it is 3-connected. We hereby denote this cycle, that is the boundary of L' , by C . We recall the definition of bridges of a cycle in a graph, which (roughly) are components of the graph minus the cycle, along with their attachments to the cycle.

► **Definition 20** ([57]). *For a subgraph H of G , a vertex of attachment of H is a vertex of H that is incident with some edge of G not belonging to H . Let J be an undirected cycle of G . We define a bridge of J in G as a subgraph B of G with the following properties:*

1. *each vertex of attachment of B is a vertex of J .*
2. *B is not a subgraph of J .*
3. *no proper subgraph of B has both the above properties.*

A bridge could also be a chord of C . We call such bridges as *trivial* bridges, and others as *non-trivial* bridges. We need the following observation for L' and C :

► **Observation 21.** *Any non-trivial bridge B of C (in the graph L'), if one exists, must have at least three vertices of attachment.*

This holds because if a non-trivial bridge of C in L' had only two vertices of attachment on C , they would form a separating pair of L .

There is a natural ordering we can give to the vertices of attachments of a bridge, that is the order in which they are encountered while travelling C , say clockwise, starting from r_0 . We now define the *span* of a bridge.

► **Definition 22.** *Let the vertices of cycle C be $\langle u_0 = r_0, u_1, u_2 \dots u_l \rangle$ in clockwise order in the given embedding. Let B be a bridge of C with vertices of attachment $u_{i_0}, u_{i_1}, \dots, u_{i_k}$, where $0 \leq i_0 < i_1 \dots i_k \leq l$. Then the set of vertices $\{u_{i_0}, u_{i_0+1}, \dots, u_{i_k}\}$ is called the span of bridge*

B with respect to r_0 (or just span of B for brevity), denoted by $\text{span}(B)$. The set of vertices $\{u_{i_k}, u_{i_k+1} \dots u_l, u_0, \dots u_{i_0}\}$ is called the complement span of bridge B (with respect to r_0). We say that bridge B has minimal span if for any other bridge B' of C , $\text{span}(B') \not\subseteq \text{span}(B)$.

Basically, the vertices of attachment of a bridge divide C into segments. The span of a bridge consists of vertices of C , excluding the ones in the segment containing r_0 , while the complement span consists of the segment containing r_0 . To disambiguate the case when r_0 itself is a vertex of attachment of a bridge B , note that the definition uses the convention of putting the segment adjacent to r_0 on the clockwise side into the span of B , and the segment adjacent to r_0 on the counter clockwise side into the complement span of B . Since G is planar and C is outer boundary, the spans of bridges form a laminar family.

► **Observation 23.** Let B_1, B_2 be two bridges of C , lying in the interior of C . Then either $\text{span}(B_1) \subset \text{span}(B_2)$, or $\text{span}(B_2) \subset \text{span}(B_1)$, or $\text{span}(B_1) \cap \text{span}(B_2) = \emptyset$.

Now we define the *second-leftmost* path of a bridge of C .

► **Definition 24.** Let B be a non-trivial bridge of cycle C , with vertices of attachment $u_{i_0}, u_{i_1}, u_{i_2} \dots u_{i_k}$ according to the ordering specified above. The second leftmost path of B starting from u_{i_0} , is the one obtained by the following procedure:

- Initialize u_{i_0} as current vertex and u_{i_0+1} as its parent.
- At any step, let the current vertex be v , and its parent be w . Let z be the next vertex occurring after w other than the vertex u_{i_1} , in the clockwise ordering of neighbours of v in the embedding. Travel to z and make it the current vertex. Make v its parent.
- Repeat the above step, until either a vertex repeats in the traversal, or a vertex of C (after the initial occurrence of u_{i_0}) is reached.

We denote this path by $P_{sl}(B, u_{i_0})$.

We next state a lemma that the second-leftmost path of a non-trivial bridge, starting at u_{i_0} , is indeed a simple path that ends at u_{i_2} because L is 3-connected.

► **Lemma 25.** Given a non-trivial bridge B of cycle C , with vertices of attachment $u_{i_0}, u_{i_1}, u_{i_2} \dots u_{i_k}$ as described above, the following hold:


1. $P_{sl}(B, u_{i_0})$ is a simple path that ends at u_{i_2} .
2. All neighbours of u_{i_1} that lie in B , also lie on $P_{sl}(B, u_{i_0})$.

Proof. We defer the proof to the full version. ◀

Since we know that $P_{sl}(B, u_{i_0})$ is a simple path starting and ending at C , it is clear that we can perform the traversal to construct it in L , by just storing the current vertex, the parent vertex, and the vertex u_{i_1} at each step. The reason we choose the *second leftmost path* in the bridge is so that after branching from the cycle C into its interior at u_{i_0} , we can emerge back at C at u_{i_2} , and then trap the maximal path in the segment $u_{i_2} \cup u_{i_0}$ of C . The existence of the unvisited vertex u_{i_1} ensures that the segment is non-empty and we can do so. In the case when the bridge B is a chord, we define its second leftmost path as just the chord itself.

We now describe the algorithm to compute maximal path in L starting from r_0 , with a small exception that we will handle subsequently.

1. Compute the bridges of C (w.r.t. L') by computing the components of $G - C$ along with their vertices of attachment to C .
2. From $r_0 = u_0$, extend the path by going clockwise till we find the first point of attachment of a bridge B of C with minimal span. Let the vertices of attachment of B , in clockwise order starting from r_0 be $u_{i_0}, u_{i_1}, u_{i_2}, \dots u_{i_k}$ (where $i_k \geq 2$). If B is a chord, we denote its vertices of attachment by u_{i_0}, u_{i_2} .



■ **Figure 5** The 3-connected graph L . The dashed edge between r_0, r_1 is a virtual edge. The boundary of $L - (r_0, r_1)$ is the cycle C . The bridge on the right side is the labelled as B , with vertices of attachment $u_{i_0}, u_{i_1}, u_{i_2}, \dots, u_{i_l}$. The maximal path is highlighted in green, starting from r_0 , going till u_{i_0} , then taking the second leftmost path in B to reach u_{i_2} , and then encircling back towards u_{i_0} , stopping at its neighbour u_{i_0+1} (In this figure, u_{i_1} is same as u_{i_0+1}).

3. From u_{i_0} , continue along the path $P_{sl}(B, u_{i_0})$ till we reach u_{i_2} (we will reach u_{i_2} by Lemma 25). From u_{i_2} , extend this path by traversing the segment $u_{i_2} \odot u_{i_0+1}$, till the vertex u_{i_0+1} . Suppose that $u_{i_0+1} \neq r_1$. Then this path is a maximal path.

To see that this path is maximal (assuming $u_{i_0+1} \neq r_1$), note that no other bridge of C can be enclosed inside $P_{sl}(B, u_{i_0})$ and the segment $u_{i_2} \odot u_{i_0}$ of C , since we chose B with a minimal span. There are two possible cases, either u_{i_0+1} has no neighbour in the interior of C in which case we are done. The only other possibility, by the above argument is that B is a non-trivial bridge and $u_{i_0+1} = u_{i_1}$ (B cannot be a chord between u_{i_0} and u_{i_0+1} as we assume input graph to be simple). In this case, by Lemma 25, its neighbours inside C are all visited in $P_{sl}(B, u_{i_0})$, and we are done.

To handle the case when $u_{i_0+1} = r_1$, we can apply the same algorithm, but some operations “reversed” (like taking the bridge with minimal *complement span*, and taking the second *rightmost* path etc). We defer this argument to the full version. Each of the steps in the algorithm like computing bridges, checking for minimal span, can be done in L using constantly many transducers and [49]. And as explained above, computing the required segments of C , and the second leftmost path, can also be done in L .

References

- 1 Alok Aggarwal and Richard J. Anderson. A random NC algorithm for depth first search. *Comb.*, 8(1):1–12, 1988. doi:10.1007/BF02122548.
- 2 Alok Aggarwal, Richard J. Anderson, and Ming-Yang Kao. Parallel depth-first search in general directed graphs. *SIAM J. Comput.*, 19(2):397–409, 1990. doi:10.1137/0219025.
- 3 Eric Allender, Archit Chauhan, and Samir Datta. Depth-first search in directed planar graphs, revisited. *Acta Informatica*, 59(4):289–319, 2022. Preliminary version appeared in MFCS 2021. doi:10.1007/S00236-022-00425-1.
- 4 Eric Allender, Archit Chauhan, Samir Datta, and Anish Mukherjee. Planarity, exclusivity, and unambiguity. *Electronic Colloquium on Computational Complexity (ECCC)*, 26:39, 2019.

- 5 Eric Allender, Samir Datta, and Sambuddha Roy. The directed planar reachability problem. In Ramaswamy Ramanujam and Sandeep Sen, editors, *FSTTCS 2005: Foundations of Software Technology and Theoretical Computer Science, 25th International Conference, Hyderabad, India, December 15-18, 2005, Proceedings*, volume 3821 of *Lecture Notes in Computer Science*, pages 238–249. Springer, 2005. doi:10.1007/11590156_19.
- 6 Eric Allender and Meena Mahajan. The complexity of planarity testing. *Information and Computation*, 189(1):117–134, 2004. doi:10.1016/j.ic.2003.09.002.
- 7 Richard Anderson. A parallel algorithm for the maximal path problem. *Combinatorica*, 7(4):315–326, 1987. doi:10.1007/BF02579320.
- 8 Richard J. Anderson and Ernst W. Mayr. Parallelism and the maximal path problem. *Inf. Process. Lett.*, 24(2):121–126, 1987. doi:10.1016/0020-0190(87)90105-0.
- 9 Sanjeev Arora and Boaz Barak. *Computational Complexity, a modern approach*. Cambridge University Press, 2009.
- 10 Nikhil Balaji. Succinct numbers, skew circuits and bounded treewidth graphs: Algorithms and complexity. In *PhD Thesis, Chennai Mathematical Institute*, 2016. URL: <https://scholar.google.co.in/scholar?oi=bibs&cluster=8516246372270295301&btnI=1&hl=en>.
- 11 Nikhil Balaji and Samir Datta. Bounded treewidth and space-efficient linear algebra. In Rahul Jain, Sanjay Jain, and Frank Stephan, editors, *Theory and Applications of Models of Computation*, pages 297–308, Cham, 2015. Springer International Publishing. doi:10.1007/978-3-319-17142-5_26.
- 12 Hans L. Bodlaender. Nc-algorithms for graphs with small treewidth. In J. van Leeuwen, editor, *Graph-Theoretic Concepts in Computer Science*, pages 1–10, Berlin, Heidelberg, 1989. Springer Berlin Heidelberg.
- 13 Chris Bourke, Raghunath Tewari, and N. V. Vinodchandran. Directed planar reachability is in unambiguous log-space. *ACM Trans. Comput. Theory*, 1(1):4:1–4:17, 2009. doi:10.1145/1490270.1490274.
- 14 Erin Wolf Chambers and David Eppstein. Flows in one-crossing-minor-free graphs. *Journal of Graph Algorithms and Applications*, 17(3):201–220, 2013. doi:10.7155/jgaa.00291.
- 15 Archit Chauhan. Some problems on planar and minor free graphs. In *PhD Thesis, Chennai Mathematical Institute*, 2024. URL: https://libarchive.cmi.ac.in/theses/archit_chauhan_cs2024.pdf.
- 16 Archit Chauhan, Samir Datta, Chetan Gupta, and Vimal Raj Sharma. The even-path problem in directed single-crossing-minor-free graphs. In *To appear in 49th International Symposium on Mathematical Foundations of Computer Science (MFCS)*, LIPIcs, 2024. doi:10.4230/LIPIcs.MFCS.2024.43.
- 17 Yijia Chen and Jörg Flum. On the ordered conjecture. In *Proceedings of the 27th Annual IEEE Symposium on Logic in Computer Science, LICS 2012, Dubrovnik, Croatia, June 25-28, 2012*, pages 225–234. IEEE Computer Society, 2012. doi:10.1109/LICS.2012.33.
- 18 Marek Cygan, Fedor V. Fomin, Lukasz Kowalik, Daniel Lokshtanov, Daniel Marx, Marcin Pilipczuk, Michal Pilipczuk, and Saket Saurabh. *Parameterized Algorithms*. Springer Publishing Company, Incorporated, 1st edition, 2015.
- 19 Samir Datta, Raghav Kulkarni, Raghunath Tewari, and N.V. Vinodchandran. Space complexity of perfect matching in bounded genus bipartite graphs. *Journal of Computer and System Sciences*, 78(3):765–779, 2012. In Commemoration of Amir Pnueli. doi:10.1016/j.jcss.2011.11.002.
- 20 Samir Datta, Nutan Limaye, Prajakta Nimbhorkar, Thomas Thierauf, and Fabian Wagner. Planar graph isomorphism is in log-space. In *Proceedings of the 24th Annual IEEE Conference on Computational Complexity (CCC)*, pages 203–214, 2009. doi:10.1109/CCC.2009.16.
- 21 Samir Datta, Nutan Limaye, Prajakta Nimbhorkar, Thomas Thierauf, and Fabian Wagner. Planar graph isomorphism is in log-space. *ACM Trans. Comput. Theory*, 14(2):8:1–8:33, 2022. doi:10.1145/3543686.

- 22 Samir Datta and Gautam Prakriya. Planarity testing revisited. In Mitsunori Ogihara and Jun Tarui, editors, *Theory and Applications of Models of Computation - 8th Annual Conference, TAMC 2011, Tokyo, Japan, May 23-25, 2011. Proceedings*, volume 6648 of *Lecture Notes in Computer Science*, pages 540–551. Springer, 2011. doi:10.1007/978-3-642-20877-5_52.
- 23 Pilar de la Torre and Clyde P. Kruskal. Fast parallel algorithms for all-sources lexicographic search and path-algebra problems. *J. Algorithms*, 19(1):1–24, 1995. doi:10.1006/jagm.1995.1025.
- 24 Reinhard Diestel. *Graph Theory*, volume 173 of *Graduate texts in mathematics*. Springer, 2016.
- 25 Michael Elberfeld, Andreas Jakoby, and Till Tantau. Logspace versions of the theorems of Bodlaender and Courcelle. In *51th Annual IEEE Symposium on Foundations of Computer Science, FOCS 2010, October 23-26, 2010, Las Vegas, Nevada, USA*, pages 143–152, 2010. doi:10.1109/FOCS.2010.21.
- 26 Michael Elberfeld and Ken-ichi Kawarabayashi. Embedding and canonizing graphs of bounded genus in logspace. In *ACM Symposium on Theory of Computing (STOC)*, pages 383–392, 2014. doi:10.1145/2591796.2591865.
- 27 David Eppstein and Vijay V. Vazirani. Nc algorithms for computing a perfect matching and a maximum flow in one-crossing-minor-free graphs. *SIAM Journal on Computing*, 50(3):1014–1033, 2021. doi:10.1137/19M1256221.
- 28 Stephen A. Fenner, Rohit Gurjar, and Thomas Thierauf. A deterministic parallel algorithm for bipartite perfect matching. *Commun. ACM*, 62(3):109–115, 2019. doi:10.1145/3306208.
- 29 Mohsen Ghaffari, Christoph Grunau, and Jiahao Qu. Nearly work-efficient parallel dfs in undirected graphs. In *Proceedings of the 35th ACM Symposium on Parallelism in Algorithms and Architectures, SPAA '23*, pages 273–283, New York, NY, USA, 2023. Association for Computing Machinery. doi:10.1145/3558481.3591094.
- 30 Martin Grohe. *Descriptive Complexity, Canonisation, and Definable Graph Structure Theory*, volume 47 of *Lecture Notes in Logic*. Cambridge University Press, 2017. doi:10.1017/9781139028868.
- 31 Chetan Gupta, Vimal Raj Sharma, and Raghunath Tewari. Reachability in $o(\log n)$ genus graphs is in unambiguous logspace. In Rolf Niedermeier and Christophe Paul, editors, *36th International Symposium on Theoretical Aspects of Computer Science, STACS 2019, March 13-16, 2019, Berlin, Germany*, volume 126 of *LIPICs*, pages 34:1–34:13. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2019. doi:10.4230/LIPICs.STACS.2019.34.
- 32 Chetan Gupta, Vimal Raj Sharma, and Raghunath Tewari. Efficient isolation of perfect matching in $o(\log n)$ genus bipartite graphs. In Javier Esparza and Daniel Král', editors, *45th International Symposium on Mathematical Foundations of Computer Science, MFCS 2020, August 24-28, 2020, Prague, Czech Republic*, volume 170 of *LIPICs*, pages 43:1–43:13. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2020. doi:10.4230/LIPICs.MFCS.2020.43.
- 33 Torben Hagerup. Planar depth-first search in $O(\log n)$ parallel time. *SIAM J. Comput.*, 19(4):678–704, June 1990. doi:10.1137/0219047.
- 34 B. Jenner and B. Kirsig. Alternierung und Logarithmischer Platz. Dissertation, Universität Hamburg, 1989.
- 35 Ming-Yang Kao. All graphs have cycle separators and planar directed depth-first search is in DNC. In John H. Reif, editor, *VLSI Algorithms and Architectures, 3rd Aegean Workshop on Computing, AWOC 88, Corfu, Greece, June 28 - July 1, 1988, Proceedings*, volume 319 of *Lecture Notes in Computer Science*, pages 53–63. Springer, 1988. doi:10.1007/BFb0040373.
- 36 Ming-Yang Kao. Linear-processor NC algorithms for planar directed graphs I: strongly connected components. *SIAM J. Comput.*, 22(3):431–459, 1993. doi:10.1137/0222032.
- 37 Ming-Yang Kao. Planar strong connectivity helps in parallel depth-first search. *SIAM J. Comput.*, 24(1):46–62, 1995. doi:10.1137/S0097539792227077.

- 38 Ming-Yang Kao and Philip N. Klein. Towards overcoming the transitive-closure bottleneck: Efficient parallel algorithms for planar digraphs. *Journal of Computer and System Sciences*, 47(3):459–500, 1993. doi:10.1016/0022-0000(93)90042-U.
- 39 Richard M. Karp and Vijaya Ramachandran. A survey of parallel algorithms for shared-memory machines. Technical Report UCB/CSD-88-408, University of California, Berkeley, March 1988. URL: <http://www2.eecs.berkeley.edu/Pubs/TechRpts/1988/5865.html>.
- 40 Richard M. Karp, Eli Upfal, and Avi Wigderson. Constructing a perfect matching in random NC. *Comb.*, 6(1):35–48, 1986. doi:10.1007/BF02579407.
- 41 Samir Khuller. Extending planar graph algorithms to $k_{3,3}$ -free graphs. *Inf. Comput.*, 84(1):13–25, 1990. doi:10.1016/0890-5401(90)90031-C.
- 42 Samir Khuller, Stephen G. Mitchell, and Vijay V. Vazirani. Processor efficient parallel algorithms for the two disjoint paths problem and for finding a kuratowski homeomorph. *SIAM J. Comput.*, 21(3):486–506, 1992. doi:10.1137/0221032.
- 43 Jan Kyncl and Tomáš Vyskocil. Logspace reduction of directed reachability for bounded genus graphs to the planar case. *ACM Trans. Comput. Theory*, 1(3):8:1–8:11, 2010. doi:10.1145/1714450.1714451.
- 44 Karl Menger. Zur allgemeinen kurventheorie. *Fundamenta Mathematicae*, 10(1):96–115, 1927. URL: <http://eudml.org/doc/211191>.
- 45 Gary L. Miller. Finding small simple cycle separators for 2-connected planar graphs. *Journal of Computer and System Sciences*, 32(3):265–279, 1986. doi:10.1016/0022-0000(86)90030-9.
- 46 Bojan Mohar and Carsten Thomassen. *Graphs on Surfaces*. Johns Hopkins series in the mathematical sciences. Johns Hopkins University Press, 2001. URL: <http://jhupbooks.press.jhu.edu/ecom/MasterServlet/GetItemDetailsHandler?iN=9780801866890&qty=1&source=2&viewMode=3&loggedIN=false&JavaScript=y>.
- 47 Ketan Mulmuley, Umesh V. Vazirani, and Vijay V. Vazirani. Matching is as easy as matrix inversion. *Comb.*, 7(1):105–113, 1987. doi:10.1007/BF02579206.
- 48 John H. Reif. Depth-first search is inherently sequential. *Inf. Process. Lett.*, 20(5):229–234, 1985. doi:10.1016/0020-0190(85)90024-9.
- 49 Omer Reingold. Undirected connectivity in log-space. *J. ACM*, 55(4):17:1–17:24, 2008. doi:10.1145/1391289.1391291.
- 50 N. Robertson and P.D. Seymour. Excluding a graph with one crossing. *Graph structure theory (Seattle, WA, 1991)*, 1993. doi:10.1090/conm/147.
- 51 Neil Robertson and P.D. Seymour. Graph minors. xx. wagner’s conjecture. *Journal of Combinatorial Theory, Series B*, 92(2):325–357, 2004. Special Issue Dedicated to Professor W.T. Tutte. doi:10.1016/j.jctb.2004.08.001.
- 52 Gregory E. Shannon. A linear-processor algorithm for depth-first search in planar graphs. *Information Processing Letters*, 29(3):119–123, 1988. doi:10.1016/0020-0190(88)90048-8.
- 53 Justin R. Smith. Parallel algorithms for depth-first searches i. planar graphs. *SIAM J. Comput.*, 15(3):814–830, 1986. doi:10.1137/0215058.
- 54 Raghunath Tewari and N. V. Vinodchandran. Green’s theorem and isolation in planar graphs. *Inf. Comput.*, 215:1–7, 2012. doi:10.1016/j.ic.2012.03.002.
- 55 Thomas Thierauf and Fabian Wagner. The isomorphism problem for planar 3-connected graphs is in unambiguous logspace. *Theory Comput. Syst.*, 47(3):655–673, 2010. doi:10.1007/s00224-009-9188-4.
- 56 Carsten Thomassen. Embeddings of graphs with no short noncontractible cycles. *Journal of Combinatorial Theory, Series B*, 48(2):155–177, 1990. doi:10.1016/0095-8956(90)90115-G.
- 57 W. T. Tutte. Separation of vertices by a circuit. *Discrete Mathematics*, 12(2):173–184, 1975. doi:10.1016/0012-365X(75)90032-1.
- 58 Klaus Von Wagner. Über eine eigenschaft der ebenen komplexe. *Mathematische Annalen*, 114:570–590, 1937. URL: <https://api.semanticscholar.org/CorpusID:123534907>.

A

 Appendix

A.1 Appendix for Section 3

A.1.1 Proof of Theorem 5

We restate the theorem here:

► **Theorem 5** ([36, 37]). *Let G be a digraph and $\{P_1, P_2\}$ be two paths that together form an α -separator of G for some $\alpha \in [\frac{1}{2}, 1)$. Given an oracle for REACH, we can in L , construct a single path P_3 that is also an α -separator of G .*

Proof. Let the two paths that together form an α separator be $P_1 = \{u_1, u_2 \dots u_p\}$ and $P_2 = \{v_1, v_2 \dots v_q\}$.

- Trim the path P_1 to $P'_1 = \{u_1, u_2 \dots u_s\}$ ($s \leq p$) such that separator property is just lost on trimming away u_{s-1} . That is, $P'_1 \cup P_2$ is an α -separator of G , but $P'_1 \setminus \{u_s\} \cup P_2$ is not. This means that in $G - (P'_1 \setminus \{u_s\} \cup P_2)$, the vertex u_s is a part of a strongly connected component of size at least αn .
- Having trimmed P_1 to P'_1 , repeat the same step and trim P_2 to $P'_2 = \{v_t, v_{t+1} \dots v_q\}$ ($t \geq 1$) such that $P'_1 \cup P'_2$ forms an α -separator, and the vertex v_t is part of some strongly connected component of size at least αn in the graph $G - (P'_1 \cup P'_2 \setminus \{v_t\})$.
- Thus in $G - (P'_1 \setminus \{u_s\} \cup P'_2 \setminus \{v_t\})$, there are two strongly connected components of size at least αn ($\alpha \geq 1/2$), containing u_s, v_t respectively. This means that they must be part of a single strongly connected component and hence there must be a path P from u_s to v_t in $G - (P'_1 \setminus \{u_s\} \cup P'_2 \setminus \{v_t\})$. Therefore we can concatenate the paths P'_1, P, P'_2 into a single path $P'_1.P.P'_2$, which is an α -separator of G .

While trimming paths in steps 1 and 2, we only need to remember the current end point of the trimmed path, which can be done in L . Other steps, like checking the size of leftover connected components and finding a path from u_s to v_t in the subgraph can be done in L given an oracle to REACH. ◀

A.2 Appendix for Section 4

We describe the steps of the algorithm for computing path separators in bounded genus digraphs here:

1. Decompose the graph into strongly connected components. If there is no strongly connected component of size larger than $2n/3$, then we are done, else let G_0 be the component larger than $2n/3$.
2. Find an embedding of G_0 using [26].
3. Find an arborescence of G_0 . To do this, first construct the path isolating weights for edges using [31, Theorem 7] so that there is a unique min weight path between any two vertices of G_0 . Then for every vertex, we find the shortest path from a vertex say r_0 to it in $\mathsf{UL} \cap \mathsf{co-UL}$ using [55] (see Section 2). This will clearly give an arborescence of G_0 rooted at r_0 .
4. By Lemma 12, we known that there exists an edge such that its fundamental cycle, is a surface non-separating cycle. To find it we can go over every edge, compute its fundamental cycle C and use the theorem of [26] to check if every weakly connected component of $G_0 - C_1$ has genus strictly lesser than that of G_0 . Let C_1 denote the surface non-separating cycle we found. This can be done in Logspace using transducers.

5. Let G_1 be the largest strongly connected component of $G_0 - C_1$. If it is more than $2n/3$ in size then we do the steps above on G_1 . Repeat until finally we have graph $G_{rem} = G - (C_1 \cup C_2 \dots C_l)$ (where $l < g$) such that either all strongly connected components of G_{rem} are smaller than $2n/3$ or all the weakly connected components are planar. In the latter case, find a $2/3$ -path separator P in the strongly connected component in G_{rem} that is of size more than $2n/3$.
6. For each cycle C_i we found at each step, let P_i, P'_i denote the two directed paths that form the cycle C_i as explained in Section 4. The paths $P_1 \cup P'_1 \dots P_l \cup P'_l \cup P$ (P if required) together form a $2/3$ -separator of G .
7. Use Corollary 6 to merge the paths $P_1, P'_1, \dots, P_l, P'_l, P$ into a single path P_s which is also a $2/3$ -path separator of G . Output P_s .

Steps 1 to 4 can be done in **L** or in **UL** \cap **co-UL**. Step 5 involves at most g iterations of steps 1 to 4 and is therefore in **UL** \cap **co-UL** as the class is closed under composition for constant number of transducers. The merging of at most $2l + 1$ paths into a single path separator is in **UL** \cap **co-UL** by Corollary 6. Therefore the entire algorithm to construct the separator is in **UL** \cap **co-UL** as it uses a constant number of transducers.