Circuit Complexity Bounds using Monoid Structure for Graph Problems

Kristoffer Arnsfelt Hansen[1], Balagopal Komarath[2], Jayalal Sarma[2], Sven Skyum[1], and Navid Talebanfard[1]

[1]Aarhus University, Denmark
[2]Indian Institute of Technology Madras, Chennai, India

December 19, 2014

Abstract

We study the complexity of several of the classical graph decision problems in the setting of bounded cutwidth and how imposing planarity affects the complexity. We show that for 2-coloring, for bipartite perfect matching, and for several variants of disjoint paths the straightforward NC^1 upper bound may be improved to $\mathsf{AC}^0[2]$, ACC^0, and AC^0 respectively for bounded planar cutwidth graphs. We obtain our upper bounds using the characterization of these circuit classes in tems of finite monoids due to Barrington and Thérien. On the other hand we show that 3-coloring and Hamilton cycle remain hard for NC^1 under projection reductions, analogous to the NP-completeness for general planar graphs. We also show that 2-coloring and (non-bipartite) perfect matching are hard under projection reductions for certain subclasses of $\mathsf{AC}^0[2]$. In particular this shows that our bounds for 2-coloring are quite close.

1 Introduction

We consider several of the classical graph decision problems, namely those of deciding existence of 2- and 3-colorings, perfect matchings, Hamiltonian cycles, and disjoint paths. For these problems we are interested in their complexity in the setting of bounded *planar cutwidth*. The *cutwidth* of a graph $G = (V, E)$ with $n = |V|$ vertices is defined in terms of linear arrangements of the vertices. A linear arrangement is simply a 1-1 map $f : V \to \{1, \ldots, n\}$, and its cutwidth is the maximum over i of the number of edges between $V_i = \{v \in V \mid f(v) \leq i\}$ and $V \setminus V_i$. The cutwidth of G is the minimum cutwidth of a linear arrangement. Similarly, if the graph G is planar we can define a notion of *planar cutwidth*. Given a linear arrangement f we consider a planar embedding where vertex v is placed at coordinate $(f(v), 0)$. The planar cutwidth of this embedding is then the maximum number of edge-crossings at a vertical line in the plane. We define the planar cutwidth as the minimum planar cutwidth of such a linear arrangement and an embedding.

All the problems we consider can be decided in NC^1 for graphs of bounded cutwidth, and they are in fact NC^1-complete under projection reductions. Imposing planarity, or more precisely considering graphs of bounded planar cutwidth, we are able to place several of the problems in smaller classes such as AC^0, $\mathsf{AC}^0[2]$, and ACC^0, while for some problems they remain NC^1-complete.

Before stating our results we review known complexity results about the graph problems without restriction on cutwidth and the consequences of imposing planarity, for comparison with our results in the bounded cutwidth setting. The 2-coloring problem is in L, as an easy consequence of Reingold's algorithm for undirected connectivity[17], whereas 3-coloring is NP-complete and remains so for planar graphs by the existence of a cross-over gadget[11]. The complexity of deciding if a graph has a perfect matching is still not known. It belongs to P, but it is an open problem if it belongs to NC. For planar graphs the problem is known to be in NC as shown by Vazirani based on work of Kasteleyn[14, 20]; for planar bipartite graphs the problem was shown to be in UL by Datta et al.[7]. The Hamiltonian cycle problem is NP-complete and as shown by Garey et al. it remains so for planar graphs[12], and Itai et al. showed it is NP-hard even for grid graphs[13].

The disjoint paths problem has numerous variations. In the general setting we are given pairs of vertices $(s_1, t_1), \ldots, (s_k, t_k)$ in a graph G, and are to decide whether disjoint paths between s_i and t_i for each i exists. Here disjoint may mean either vertex-disjoint or edge-disjoint, but either variant is reducible to the other. We shall consider only the case of constant k. When G is an undirected graph a polynomial time algorithm was given by Robertson and Seymour[18], as a result arising from their seminal work on graph minors. When G is a directed graph the problem is NP-complete already for $k = 2$ as shown by Fortune et al.[10]. On the other hand, when G is planar Schrijver[19] gave a polynomial time algorithm for the vertex-disjoint paths problem. The reduction between the vertex-disjoint and the edge-disjoint versions of the problem does not preserve planarity, and it is an open problem the edge-disjoint paths problem in planar directed graphs is NP-complete or solvable in polynomial time[6]. It can however be solved in polynomial time for (not necessarily planar) directed acyclic graphs[10].

1.1 Results and techniques

A convenient way to obtain an NC^1 upper bound is through monadic second order (MSO) logic. Elberfeld et al.[8] showed that MSO-definable problems can be decided in NC^1 when restricted to input structures of bounded treewidth, when a tree decomposition of bounded width is supplied in the so-called term representation. We shall not formally define the treewidth of a graph, but we will note that the treewidth of a graph is bounded from above by the cutwidth of the graph[5]. Furthermore, given as input a linear arrangement of bounded cutwidth k, a tree decomposition of tree width k can be constructed by an AC^0 circuit. Thus we have the following meta-theorem as an easy consequence.

Theorem 1. *Any graph property definable in monadic second order logic with quantification over sets of vertices and edges can be decided by* NC^1 *circuits on graphs of bounded cutwidth if a linear arrangement of bounded cutwidth is supplied as auxiliary input.*

Actually to obtain this we may proceed more directly avoiding the challenges dealt with in [8], since the tree decomposition computed above is actually a path decomposition. We may thus also use standard finite automata rather than tree automata.

All the graph properties we consider can easily be expressed in monadic second order logic, thereby establishing NC^1 upper bounds. We can show that all these problems are in fact also hard for NC^1 under projection reductions. This is based on Barrington's characterization of NC^1 in terms of bounded width permutation branching programs[2].

We first discuss the general technique behind our upper bounds that improve upon the generic NC^1 bound. Namely our upper bounds are based on reducing to *word problems* on appropriately defined finite monoids. By results of Barrington and Thérien, we then directly get circuit upper bounds depending on the group structure of the given monoid. For general graphs the improved complexity bounds obtained when imposing planarity are obtained by very different algorithms. In our setting of constant cutwidth, when imposing planarity we instead obtain the improvements in a uniform way by obtaining an *algebraic* understanding of the respective problems. The general idea is as follows. We consider grid-planar graphs (defined later) of a fixed width w for which we want to decide a certain graph property, and we may view these as a free semigroup under concatenation. We then define an appropriate finite monoid $\mathcal{M}$. For each grid-planar graph G we associate a monoid element $G^{\mathcal{M}}$. In the simplest setting we will be able to determine if the graph property under consideration holds for the graph G directly from the monoid element $G^{\mathcal{M}}$. We will also have defined the elements of $\mathcal{M}$ and the monoid operation in such a way that the map $G \mapsto G^{\mathcal{M}}$ is a homomorphism. What then remains is to analyze the groups inside $\mathcal{M}$. For the disjoint paths problem we show that all groups are trivial, and this gives AC^0 circuits. For 2-coloring we characterize the groups as being isomorphic to groups of the form $\mathbb{Z}_2^l$, and this gives $\mathsf{AC}^0[2]$ circuits. For perfect matching in bipartite graphs we are not able to fully analyze the groups of the corresponding monoid. We are however able to rule out groups of order 2, and thus by the celebrated Feit-Thompson theorem all remaining groups must be solvable, and this gives ACC^0 circuits.

When considering the graph properties for graphs of bounded planar cutwidth we supply as additional input the corresponding embedding of bounded cutwidth of the graph. But before dealing with this issue, we consider special classes of such graphs where such an embedding is implicit. We consider a grid $\Lambda = \{1, \ldots, l\} \times \{1, \ldots, w\}$ of *width* w and *length* l. A *grid graph* $G = (V, E)$ of width w and length l is a graph where $V \subseteq \Lambda$ and all edges are of Euclidean length 1. We think of the vertices with the same first coordinate to be in the same *layer*. A grid graph with (planar) diagonals allows edges of Euclidean length < 2, but no crossing edges. We relax these requirements further, defining the class of constant width *grid-planar* graphs. A grid-planar graph $G = (V, E)$ of width w and length l is a graph where $V \subseteq \Lambda$ and if two vertices (a, b) and (c, d) are connected by an edge, then $|a - b| \leq 1$ and the edge is fully contained in the region $[a-1, a] \times [1, w]$ or the region $[a, a+1] \times [1, w]$.

If we consider bipartite grid-planar graphs we assume that the bipartition is defined by the parities of the sums of coordinates of each vertex. All our lower bounds hold for grid graphs or grid graphs with diagonals, and all our circuit upper bounds hold for grid-planar graphs.

Just as 3-coloring and Hamiltonian cycle remain NP-complete for planar graphs, 3-coloring remains hard for NC^1 on constant width grid graphs with diagonals and Hamiltonian cycle remains hard for NC^1 on constant width grid graphs. We show that 2-coloring on constant width grid-planar graphs is in $\mathsf{AC}^0[2]$. This is complemented by an $\mathsf{AND} \circ \mathsf{XOR} \circ \mathsf{AC}^0$ lower bound for grid graphs with diagonals. This lower bound is in some sense not far from the $\mathsf{AC}^0[2]$ upper bound. Namely by the approach of Razborov[16] we have that quasipolynomial size randomized XOR $\circ$ AND is equal to quasipolynomial $\mathsf{AC}^0[2]$. Furthermore Allender and Hertrampf[1] show that in fact quasipolynomial size $\mathsf{AND} \circ \mathsf{OR} \circ \mathsf{XOR} \circ \mathsf{AND}$ is equal to quasipolynomial size $\mathsf{AC}^0[2]$.

We show that perfect matching is in ACC^0 for bipartite grid-planar graphs, and we have an AC^0 lower bound. For non-bipartite grid-planar graphs we have a $\mathsf{AND} \circ \mathsf{OR} \circ \mathsf{XOR} \circ \mathsf{AND}$ lower bound. For the disjoint paths problem in constant width grid-planar graphs we give AC^0 upper bounds for the following 3 settings: (1) node-disjoint paths in directed graphs. (2) edge-disjoint paths in upward planar graphs. (3) edge-disjoint paths in undirected graphs. We leave open the case of edge-disjoint paths in directed graphs. For all the settings we have an AC^0 lower bound. All these results are summarized in Figure 1.

Problem	Upper bound	Lower bound (projections)
2-coloring	$\mathsf{AC}^0[2]$	$\mathsf{AND} \circ \mathsf{XOR} \circ \mathsf{AC}^0$
3-coloring	NC^1	NC^1
Bipartite perfect matching	ACC^0	AC^0
Perfect matching	NC^1	$\mathsf{AND} \circ \mathsf{OR} \circ \mathsf{XOR} \circ \mathsf{AC}^0$
Hamiltonian cycle	NC^1	NC^1
Disjoint paths variants	AC^0	AC^0

Figure 1: Complexity of problems on constant width grid-planar graphs

We shall now discuss extending the upper bounds above from constant width grid-planar graphs to the larger classes of graphs of bounded planar cutwidth. Whereas the embedding was implicitly given for grid-planar graphs, for graphs of bounded planar cutwidth we will supply a representation of the embedding in addition to the linear arrangement of the vertices. A simple way to represent both the linear arrangement and the planar embedding of a graph $G = (V, E)$ of bounded planar cutwidth is to provide instead a grid-planar graph $G' = (V', E')$, where $V \subseteq V'$, where the vertices V are placed on a horizontal line and the vertices $V' \setminus V$ are dummy vertices describing the embedding of the edges. When given this representation as input, our upper bounds are easily adapted. Namely, for the disjoint paths problems an edge can be replaced by a path, and we may simply promote the dummy vertices to regular vertices. For 2-coloring and perfect matching an edge can be replaced by a path of odd length, but this can be done by an AC^0 circuit making locally use of the coordinates of vertices. Namely, we can just ensure that the path alternates between vertices of the implicit bipartition, except possibly at the end.

2 Preliminaries

Boolean circuits We give here standard definitions of the Boolean functions and circuit classes we consider. As is usual, when considering a Boolean function $f : \{0,1\}^n \to \{0,1\}$, unless otherwise specified we always have a family of such functions in mind, one for each input length. AC^0 is the class of polynomial size constant depth circuits built from unbounded fanin AND and OR gates. $\mathsf{AC}^0[m]$ allows in addition the function MOD_m given by $\mathrm{MOD}_m(x_1, \dots, x_k) = 1$ if and only if $\sum_{i=1}^{k} x_i \not\equiv 0 \pmod{m}$. We shall also denote the function MOD_2 by XOR. The union of $\mathsf{AC}^0[m]$ for all m is the class ACC^0. NC^1 is the class of polynomial size circuits of depth $O(\log n)$ built from fanin 2 AND and OR gates.

A class of Boolean functions immediately defines a class of Boolean circuits as families of single gate circuits. Given two classes of circuits $\mathcal{C}_1$ and $\mathcal{C}_2$ we denote by $\mathcal{C}_1 \circ \mathcal{C}_2$ the class of circuits consisting of circuits from $\mathcal{C}_1$ that is fed as inputs the output of circuits from $\mathcal{C}_2$. For instance, $\mathsf{AND} \circ \mathsf{XOR} \circ \mathsf{AC}^0$ is the class of polynomial size constant depth circuits that has an AND gate at the output, followed by XOR gates that in turn take as inputs the output of AC^0 circuits.

Semigroups, monoids and programs A semigroup is a set $\mathcal{S}$ with an associative binary operation. A monoid $\mathcal{M}$ is a semigroup with a two-sided identity. A subset $\mathcal{G}$ of $\mathcal{M}$ is a group in $\mathcal{M}$ if it is a group with respect to the operation of $\mathcal{M}$. We also say that $\mathcal{M}$ contains $\mathcal{G}$. A monoid is aperiodic if every group it contains is trivial; it is solvable if every group it contains is solvable. A monoid which is not solvable is called unsolvable.

We consider the *program over monoid* [4] formalism for computing Boolean functions. Let M be a monoid and n an input length. An *instruction* is a triple $\langle j, a_0, a_1 \rangle$, where $j \in [n]$ and $a_0, a_1 \in M$. A program over M is a pair (P, A) where $A \subset M$ is the accepting set and $P = (I_1, \dots, I_\ell)$ is a list of instructions. The length of the program is ℓ. Let $x \in \{0,1\}^n$. The output $I(x)$ of an instruction $I = \langle j, a_0, a_1 \rangle$ is a_{x_j}. The (Boolean) output of the program is is 1 if and only if $\prod_{i=1}^{\ell} I_i(x) \in A$. As with circuits we consider families of programs, one for each input length.

Barrington and Thérien[4] showed that several circuit classes are exactly captured by programs over finite monoids of polynomial length.

Theorem 2 (Barrington and Thérien). *Let $L \subseteq \{0,1\}^n$.*

- *L is in AC^0 if and only if L is computed by a polynomial length programs over a finite aperiodic monoid.*
- *L is in $\mathsf{AC}^0[m]$ if and only if L is computed by a polynomial length program over a finite solvable monoid in which all groups have orders dividing a power of m[1].*

[1] This equivalence is not stated explicitly in [4], but follows from the given proof.

- *L is in* ACC^0 *if and only if L is computed by a polynomial length program over a finite solvable monoid.*
- *L is in* NC^1 *if and only if L is computed by a polynomial length program over a finite unsolvable monoid.*

We shall use only one direction of this characterization, and for this reason it is convenient to reformulate as follows. Let $\mathcal{M}$ be a finite monoid. The *word problem* over $\mathcal{M}$ is to compute the product $x_1 \cdots x_m$ when given as input $x_1, \ldots, x_m \in \mathcal{M}$. When $\mathcal{M}$ is aperiodic the word problem is in AC^0, when $\mathcal{M}$ is solvable and all groups in $\mathcal{M}$ have orders dividing a power of m the word problem is in $\mathsf{AC}^0[m]$, when $\mathcal{M}$ is solvable the word problem is in ACC^0, and we always have the word problem is in NC^1.

3 Upper bounds

We first state a geometric lemma that we shall make use of in our results about bipartite matching and disjoint paths. Consider a piecewise smooth infinite simple curve C such that C is contained entirely in the strip $\{(x, y) \mid 1 \leq y \leq w\}$. We say that C is periodic with period p if the horizontally shifted curve $C + (p, 0)$ coincides with C.

Lemma 3. *Let C be a curve that is periodic with period p and let $C' = C + (q, 0)$ be a horizontal shift of the curve C. Then C and C' intersect.*

Proof. Map the region $R = \{(x, y) \mid 0 \leq x \leq p, 1 \leq y \leq w\}$ into the plane by the map $\phi(x, y) = (y \cos(2\pi x/p), y \sin(2\pi x/p))$. Then $\phi(C)$ and $\phi(C')$ are closed simple curves in the plane both containing the origin. By the Jordan curve theorem each of these curves divide the plane into an inside set and an outside set. If they do not intersect then either $\phi(C)$ encloses $\phi(C')$ or $\phi(C')$ encloses all of $\phi(C)$. In particular this means that the two curves enclose sets of different areas. However $\phi(C')$ is just a rotation of $\phi(C)$ around the origin, and must in particular enclose the exact same area as $\phi(C)$. We conclude the curves intersect. □

3.1 2-coloring

We prove here our upper bound for 2-coloring.

Theorem 4. *Testing whether a given grid-planar graph is 2-colorable can be done in* $\mathsf{AC}^0[2]$.

We prove this result by reducing 2-coloring to the word problem over a finite monoid $\mathcal{M}$. We then show that all groups in $\mathcal{M}$ are solvable and of order a power of 2. By Theorem 2 this gives us our $\mathsf{AC}^0[2]$ upper bound.

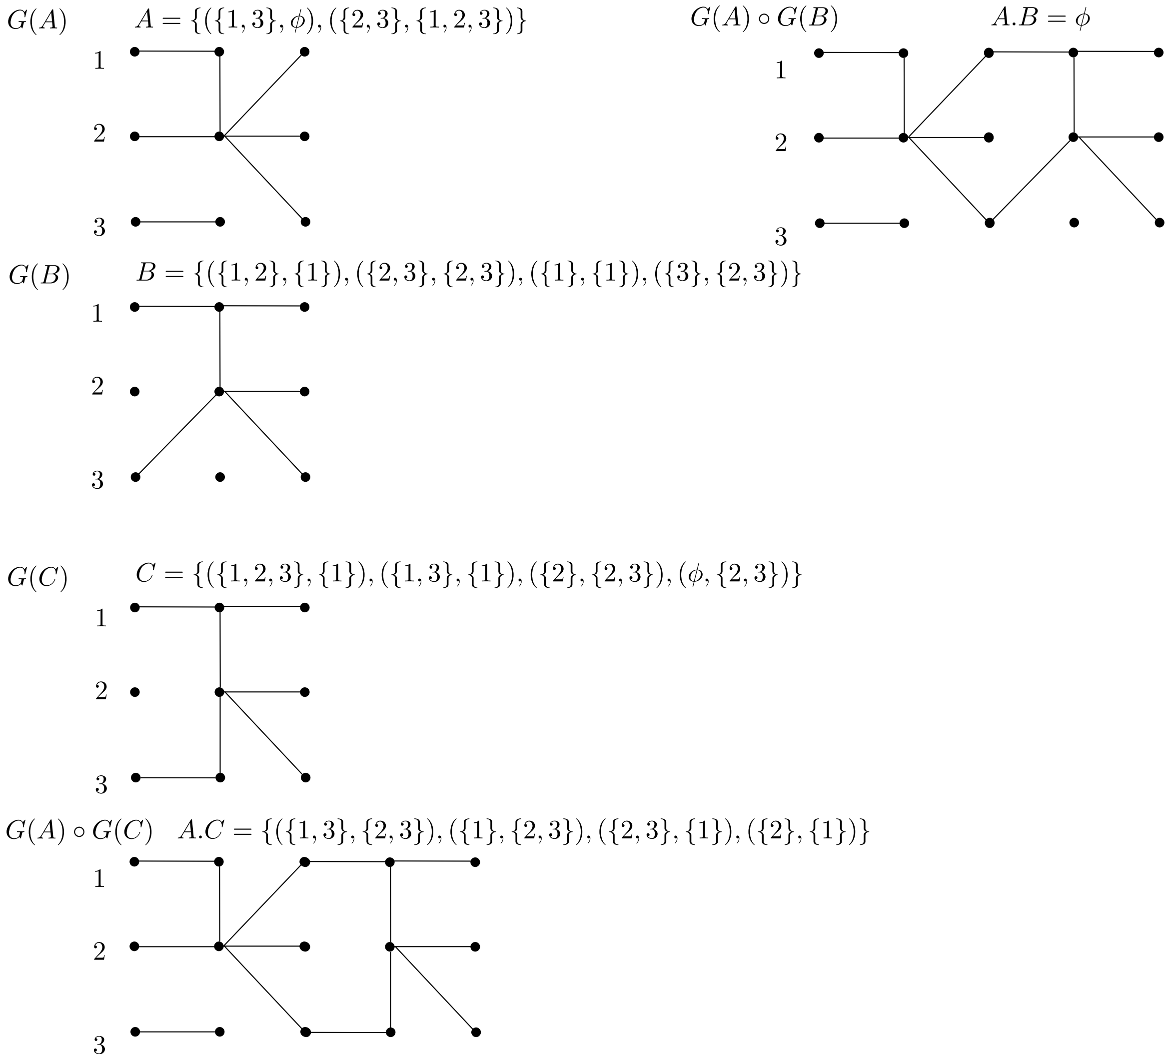

Figure 2: Monoid for 2-coloring

Reduction to a Monoid Word Problem: A grid-planar graph G gives rise to a binary relation $\mathcal{R}(G) \subseteq 2^{\{1,\ldots,w\}} \times 2^{\{1,\ldots,w\}}$. Here $\{1, \ldots, w\}$ are the numbering of the vertices in the layers of the graph. We have that $(S, T) \in \mathcal{R}(G)$ if and only if there is a 2-coloring of G such that the vertices in the first layer colored 1 is the set S and the vertices in the last layer colored 1 is the set T. Some examples are shown in Figure 2. Let $\mathcal{M}$ be the monoid of all such relations under normal composition of relations. Let $G \circ H$ denote the concatenation of the graphs G and H.

Lemma 5. $\mathcal{R}(G \circ H) = \mathcal{R}(G)\mathcal{R}(H)$

Proof. $\subseteq$: Let $(S, T) \in \mathcal{R}(G \circ H)$. Consider any 2-coloring of $G \circ H$ that witnesses this. Under this coloring all vertices in S in the first layer and all vertices in T in the last layer are assigned color 1. Let U be the set of vertices that get color 1 on the layer that G and H meet. Then obviously $(S, U) \in \mathcal{R}(G)$ and $(U, T) \in \mathcal{R}(H)$ which means $(S, T) \in \mathcal{R}(G)\mathcal{R}(H)$.

$\supseteq$: Let $(S, T) \in \mathcal{R}(G)\mathcal{R}(H)$. Then there exists U such that $(S, U) \in \mathcal{R}(G)$ and $(U, T) \in \mathcal{R}(H)$. Let σ and π be 2-colorings of G and H, respectively, that witness this. Then we can get a 2-coloring of $G \circ H$ by $\sigma\pi$, i.e., by coloring G using σ and then extending it via π on H. □

The proof of the upper bound is now completed by the following result.

Proposition 6. *Every group $\mathcal{G} \subseteq \mathcal{M}$ is isomorphic to $\mathbb{Z}_2^\ell$ for some ℓ.*

Proof. For a graph G, let us identify the nodes in the first layer with the set $\{1, \dots, w\}$ and the nodes in the last layer with the set $\{1', \dots, w'\}$.

The observation that makes the proof possible is the following: Suppose G and H are both 2-colorable graphs and $\mathcal{R}(G) = \mathcal{R}(H)$. Then for any $u, v \in \{1, \dots, w\} \cup \{1', \dots, w'\}$ we have that u and v are connected in G if and only if u and v are connected in H. Furthermore if u and v are connected in G by an odd (even) length path then u and v are connected in H by an odd (even) length path.

Assume that G is 2-colorable. Then every connected component of G can be 2-colored in exactly two different ways. This means that $\mathcal{R}(G)$ can be reconstructed from only the following information about G: Which nodes from $\{1, \dots, w\} \cup \{1', \dots, w'\}$ are connected in addition to a single 2-coloring of those vertices.

Let $\mathcal{G} \subseteq \mathcal{M}$ be a nontrivial group with identity E. For any element A in $\mathcal{G}$ fix a grid-planar graph $G(A)$ such that $\mathcal{R}(G(A)) = A$. Note that each such $G(A)$ is 2-colorable. Otherwise A is the empty relation, and since that behaves as a zero element in $\mathcal{M}$, the group $\mathcal{G}$ would be the trivial group consisting just of the empty relation (since each element of $\mathcal{G}$ must have an inverse).

Let $(i \sim j) \in G$ denote that i and j are connected via a path in G.

Claim 1. Let $A, B \in \mathcal{G}$. Then $(i \sim j) \in G(A)$ if and only if $(i \sim j) \in G(B)$, and $(i' \sim j') \in G(A)$ if and only if $(i' \sim j') \in G(B)$.

Proof. It is enough to prove the claim for the case when B is just the identity element E. Assume $(i \sim j) \notin G(A)$. Since $EA = A$ and that $\mathcal{R}(G(E) \circ G(A)) = EA = A$ by Lemma 5, there cannot be a path between i and j in $G(E) \circ G(A)$ since otherwise the color of i and j will depend on each other and we know that this is not the case since $(i \sim j) \notin G(A)$. Therefore $(i \sim j) \notin G(E) \circ G(A)$ which in particular implies $(i \sim j) \notin G(E)$. For the other direction assume that $(i \sim j) \notin G(E)$. We have $AA^{-1} = E$. Using Lemma 5 we get $\mathcal{R}(G(A) \circ G(A^{-1})) = AA^{-1} = E$. This implies that there is no path between i and j in $G(A) \circ G(A^{-1})$ since otherwise the colors of i and j will depend on each other in $G(E)$ which we know is not the case. Therefore $(i, j) \notin G(A) \circ G(A^{-1})$ and hence $(i, j) \notin G(A)$. This shows that $(i, j) \in G(E)$ if and only if $(i, j) \in E(A)$. To show that $(i' \sim j') \in G(A)$ if and only if $(i' \sim j') \in G(E)$ we consider equations $AE = A$ and $A^{-1}A = E$ and use a similar argument as above. □

Let $A \in \mathcal{G}$ and consider the graph $G(A)$. For any set $S \subseteq G(A)$ let $\overleftarrow{V}(S) = S \cap \{1, \ldots, w\}$ and $\overrightarrow{V}(S) = S \cap \{1', \ldots, w'\}$. We define $L(A)$ to be the set of all connected components C in $G(A)$ such that $\overleftarrow{V}(C) \neq \emptyset$ and $\overrightarrow{V}(C) = \emptyset$. Similarly let $R(A)$ denote the set of all connected components C such that $\overleftarrow{V}(C) = \emptyset$ and $\overrightarrow{V}(C) \neq \emptyset$. We now let $V_L(A) = \{\overleftarrow{V}(C) : C \in L(A)\}$ and $V_R(A) = \{\overrightarrow{V}(C) : C \in R(A)\}$. Define $M(A)$ to be the set of connected components that are neither in $L(A)$ nor in $R(A)$ and have vertices on both sides of $G(A)$. We then define $V_L^M(A) = \{\overleftarrow{V}(C) : C \in M(A)\}$ and $V_R^M(A) = \{\overrightarrow{V}(C) : C \in M(A)\}$. ve

Claim 2. The following properties hold.

(i) $V_L(A) = V_L(E)$ and $V_R(A) = V_R(E)$. Furthermore, for any pair of i and j that are in the same component in $L(A)$, the lengths of all paths between i and j in $G(A)$ have the same parity and that is the same as in $G(E)$. Similarly for every pair i' and j' that are in the same connected component in $R(A)$, the length of all paths between i' and j' are of the same parity and that is the same as in $G(E)$.

(ii) $V_L^M(A) = V_L^M(E)$ and $V_R^M(A) = V_R^M(E)$.

Proof. All these follow straightforwardly from $A^{-1}A = AA^{-1} = E$ and $EA = AE = A$.

(i) Consider a node i in $G(E)$ which appears in some component in $L(E)$. Note that $AA^{-1} = E$. Since the color of i in $G(E)$ does not depend on any node on the right side of $G(E)$ the same should hold for $G(A) \circ G(A^{-1})$. By Claim 1 we know that for any j we have $(i \sim j) \in G(A)$ if and only if $(i \sim j) \in G(E)$. Therefore $V_L(E) = V_L(A)$. Considering $A^{-1}A = E$ and using a similar argument we can show that $V_R(E) = V_R(A)$. The parity of path lengths are preserved because since $G(A) \circ G(A^{-1})$ and $G(E)$ admit the same 2-colorings.

(ii) Let i and j be nodes in some set $S \in V_L^M(A)$. Consider $H = G(E) \circ G(A)$. Then by Claim 1 we know that i and j are connected $G(E)$. But they should also be connected to some node on the right end layer of H, since H and $G(A)$ admit the same colorings on boundaries. This means that i and j are both in some set $T \in V_L^M(E)$. Conversely let i and j be nodes in some set $S \in V_L^M(E)$. This time we let $H = G(A) \circ G(A^{-1})$. Again by Claim 1, i and j are connected in $G(A)$, but they are also connected to some node on the right end layer of $G(A^{-1})$. This means that they are also connected to some nodes on the right end layer of $G(A)$ and thus they belong to some set $T \in V_L^M(A)$. This shows that $V_L^M(A) = V_L^M(E)$. We can prove $V_R^M(A) = V_R^M(E)$ similarly.

□

For each component in $M(A)$ we pick two representatives, one from each side. We pick the left representatives $i_1 < \ldots < i_m$ arbitrarily. But for the right representatives if i_k is connected to i'_k then we pick i'_k as the representative of the k'th component, otherwise we pick an arbitrary node in the component. Let the right representatives be $j'_1 < \ldots < j'_m$. We map the left representative of a component to its right representative. Since $G(A)$ is planar, for every k we have that i_k is mapped to j'_k. This means that we can rename the components in $M(A)$ by $C_1, \ldots, C_m$ such that all vertices in C_i appear after all vertices in C_{i-1}.

Furthermore we know by above claim that in a group, these components are the same on the boundaries of the graph of each group element. For any $A \in \mathcal{G}$ and any $1 \leq k \leq m$ let π_k^A be the parity of the length of all paths between i_k and j'_k in $G(A)$. We show that there exists a sequence $\epsilon_1, \ldots, \epsilon_m \in \{0,1\}^m$ such that for any $A, B \in \mathcal{G}$ and all $1 \leq k \leq m$, the parity of the paths between i_k and j'_k in $G(A) \circ G(B)$ is given by $\pi_k^A \oplus \pi_k^B \oplus \epsilon_k$. To see this consider the graph $G(A) \circ G(B)$ and rename the i_k and j'_k on the side where $G(A)$ and $G(B)$ meet as $i^{(1)}$ and $i^{(2)}$ ($i^{(1)} = i_k$ and $i^{(2)} = j'_k$). If $j'_k = i'_k$ we set $\epsilon_k = 0$. This clearly satisfies the desired property, since to get from i_k on the left layer of $G(A)$ to i'_k on the right layer of $G(B)$ we can first go to i'_k on the right layer of $G(A)$ and then to i'_k on the right layer of $G(A)$. Any such path has clearly parity $\pi_k^A + \pi_k^B$. If $j'_k \neq i'_k$ we note that the parity between $i^{(1)}$ and $i^{(2)}$ is exactly the same as in $G(E) \circ G(E)$ by Claim 2. We denote this by ϵ_j. Now to color $G(A) \circ G(B)$ if we use color 0 on i_j then we are forced to use color π_j^A on $i^{(2)}$, and hence $\pi_j^A \oplus \epsilon_j$ on $i^{(1)}$ and finally we should use $\pi_j^A \oplus \epsilon_j \oplus \pi_j^B$ on i'_j. This means that the parity between i_j and i'_j is $\pi_j^A \oplus \pi_j^B \oplus \epsilon_j$ as claimed.

We define a group $\mathbb{Z}_2^{(\epsilon_1,\ldots,\epsilon_m)}$ as follows. The elements are just the same as $\mathbb{Z}_2^m$, and the group operation is defined as $\mathbb{Z}_2^m$ but then adding the vector $(\epsilon_1, \ldots, \epsilon_m)$ to the result. It is clear that $\mathbb{Z}_2^{(\epsilon_1,\ldots,\epsilon_m)}$ and $\mathbb{Z}_2^m$ are isomorphic. The above argument shows that $\mathcal{G}$ is isomorphic to $\mathbb{Z}_2^{(\epsilon_1,\ldots,\epsilon_m)}$ and hence to $\mathbb{Z}_2^m$ □

3.2 Bipartite matching

We prove here our upper bound for bipartite matching.

Theorem 7. *Given a bipartite grid-planar graph G, we can decide whether G has a perfect matching in* ACC^0.

Reduction to a Monoid Word Problem For each grid-planar graph G of odd length ℓ that has no vertical edges in the rightmost layer, we define the corresponding monoid element $G^{\mathcal{M}}$ as the triple (X, Y, R) where $X \subseteq [w]$ is the set of vertices in the leftmost layer of G, $Y \subseteq [w]$ is the set of vertices in the rightmost layer of G and $R \subseteq 2^X \times 2^Y$ is a binary relation such that for any $X_1 \subseteq X$, $X_2 \subseteq Y$ we have $(X_1, X_2) \in R$ if and only if G has a matching that matches all vertices in G except $\overline{X_1}$ in the leftmost layer and X_2 in the rightmost layer. The monoid product is defined as $(X_1, X_2, R)(X_3, X_4, S) = (X_1, X_4, R \circ S)$ when $X_2 = X_3$ and $\circ$ is the usual composition of binary relations. When $X_2 \neq X_3$, we define the product to be an element 0 for which $0x = x0 = 0$ for any x in the monoid. Now define the monoid $\mathcal{M} = \{G^{\mathcal{M}} : G \text{ is an odd length bipartite grid-planar graph}\} \cup \{0\} \cup \{1\}$, where 1 is an added identity. It is easy to see that the monoid operation described corresponds to concatenation of graphs (by merging the vertices in the rightmost layer of first graph with the vertices in the leftmost layer of the second graph). So a perfect matching exists in G

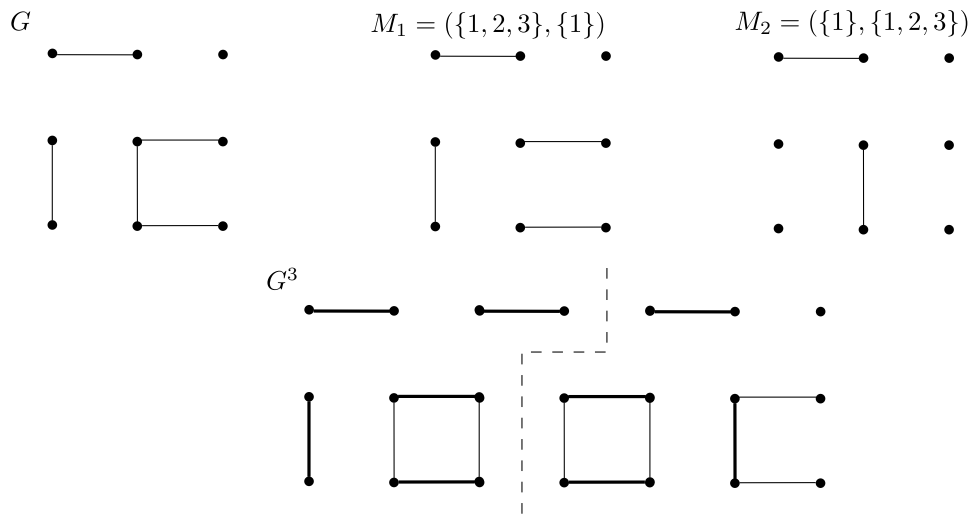

Figure 3: Monoid for Bipartite Perfect Matching

if and only if $G^{\mathcal{M}} = (X_1, X_2, R)$ and R contains the element $(X_1, \overline{X_2})$. Some examples for bipartite grid graphs and their corresponding elements in the monoid are shown in Figure 3.

We begin by considering an arbitrary group $\mathcal{G} \subset \mathcal{M}$, such that $\mathcal{G} \neq \{0\}$. First, observe that for any two elements (X_1, X_2, R) and (X_3, X_4, S) in the group $X_1 = X_2 = X_3 = X_4$ as $0 \notin \mathcal{G}$. So we can identify any element (X_1, X_2, R) of the group by simply using R. Let $o_{\mathcal{G}}(R)$ denote the order of the element R in the group $\mathcal{G}$. Let E be the identity element in $\mathcal{G}$. Now suppose an $R \in \mathcal{G}$ such that $R \neq E$ exists. The proof outline is as follows. First we relate $o_{\mathcal{G}}(R)$ to length of *primary cycles* in R. Then we show that $o_{\mathcal{G}}(R) \neq 2$ for arbitrary $R \in \mathcal{G}$. We then use Proposition 8 and Theorem 9 to conclude that $\mathcal{G}$ is solvable.

Proposition 8. *If $\mathcal{G}$ is a finite group of order $2k$ for some $k \geq 1$, then there exists $a \in \mathcal{G}$ such that $a \neq e$ and $a^2 = e$.*

Proof. Let $a_1, \ldots, a_{2k-1}$ be the non-identity elements in $\mathcal{G}$. Pair each a_i with its inverse a_j. There will be at least one a_i such that $a_i = {a_i}^{-1}$. So ${a_i}^2 = e$ where e is the identity element in $\mathcal{G}$. □

Theorem 9 (Feit-Thompson [9])**.** *Every group of odd order is solvable.*

Definition 10 (Primary Cycle)**.** A cycle $C = X_1 \to \ldots X_n \to X_1$ in the relation digraph of R is called a *primary cycle* if and only if

1. The cycle C is the smallest cycle in $R[C]$ (The subgraph induced by C).

2. The relation E does not contain (X_i, X_j) where $i \neq j$.

Lemma 11. *Every $R \in \mathcal{G}$ where $R \neq E$ must contain a primary cycle of length at least 2.*

Proof. Let us observe the structure of relation digraph representing R. It is clear that R must contain some cycle (not just self-loops). Suppose it does not, then let ℓ be the length of the smallest simple path (in edges) in R (not taking self-loops). Then any edge in $R^{\ell+1}$ must be obtained by taking a self-loop at least once. So if we take this self-loop one more time, we can obtain this edge in $R^{\ell+2}$. Similarly, for any edge in $R^{\ell+2}$, we must take some self-loop at least twice or it must take at least two self-loops. Therefore, taking one of these self-loops one less time gives this edge in $R^{\ell+1}$. So $R^{\ell+1} = R^{\ell+2}$, which implies $R \notin \mathcal{G}$.

Now we will argue that R must have at least one primary cycle of length at least 2. Notice that if R has some cycle then it must have some induced cycles. If all such cycles have self-loops (i.e., for any cycle in the graph there is a chord or a self-loop on one of its vertices), then by an argument similar to the one in the previous paragraph $R^k = R^{k+1}$ for some k which is a contradiction. So R must contain some induced cycle of length at least 2. Now if for this cycle C, E contains (X_i, X_j) for $i \neq j$ and $X_i, X_j \in C$, then using $RE = ER = R$ we conclude that C has a chord or a self-loop at one of the vertices and hence it is not an induced cycle. Now since $R \in \mathcal{G}$, we have for some k that $R^k = E$. Then the relation E must have self-loops at all vertices in C. Therefore C is a primary cycle. □

Now that we've established that any non-identity element must have at least one primary cycle of length at least two, we present the following claim relating the order of an element to its primary cycles.

Claim 3. $o_{\mathcal{G}}(R)$ is a common multiple of the lengths of primary cycles in R.

Proof. Suppose there exists a primary cycle C of length n in R such that n does not divide $m = o_{\mathcal{G}}(R)$. Then $R^m = E$ contains an edge (X_i, X_j) for some $X_i, X_j \in C$, $i \neq j$ which contradicts the assumption that C is a primary cycle. □

Claim 4. The monoid $\mathcal{M}$ is solvable.

Proof. We will prove that for any R, $o_{\mathcal{G}}(R) \neq 2$. Suppose $o_{\mathcal{G}}(R) = 2$, then using Lemma 11 and Claim 3 we can conclude that R has a primary cycle $C = X_1 \to X_2 \to X_1$ of length 2. Consider a graph G defining R, and let M_1 be a matching in G corresponding to $(X_1, X_2) \in R$ and let M_2 be a matching in G corresponding to $(X_2, X_1) \in R$. Consider now the graph $S = M_1 \cup M_2$. The graph S^n is obtained by concatenating n copies of S. We note that for any odd (even) n, the graph S^n is a union of two matchings. The matching M obtained by the concatenation of matchings $M_1 M_2 \dots$ and the matching N obtained by the concatenation of matchings $M_2 M_1 \dots$.

We label the vertices on the left side on the i^{th} copy of S as $1_{(i)}, \dots, k_{(i)}$. The rightmost vertices in S^n are labeled $1_{(n+1)}, \dots, k_{(n+1)}$.

A path in S^n is called a *blocking path* if it connects some vertex in the leftmost layer to some vertex in the rightmost layer.

Claim 5. For any n, the graph S^n must have a blocking path.

Proof. Suppose S^n does not have any blocking path. Assume wlog n is even, then consider the set V_L of all vertices in S^n that are reachable from some vertex in the left end and the set V_R of all vertices in S^n that are reachable from some vertex in the right end. Put any remaining vertices in the set V_L. Since there is no blocking path V_L and V_R are disjoint. Now we can obtain a matching (X_1, X_2) in $R^n = E$ by using the matching M_1 on the vertices in V_L and using the matching M_2 on V_R. This is a contradiction. □

We say that a path P crosses a boundary in S^n if it has two consecutive edges e_1 and e_2 such that they belong to different copies of S in S^n. Note that e_1 and e_2 must belong to the same matching M_1 or M_2. If they do not, the vertex common to those edges must be in $X_1 \cap \overline{X_1}$ or $X_2 \cap \overline{X_2}$.

Claim 6. For any n, the graph S^n cannot have a path from $v_{(i)}$ to $v_{(i+1)}$ for any i and v.

Proof. To simplify the proof, for a graph corresponding to a given monoid element we attach length 2 horizontal paths to the vertices in the left and right side through two new layers. Notice that this does not change the monoid element since the vertices in the graph corresponding to the monoid element which were originally matched inside the graph remains matched inside the graph itself and vice versa.

Suppose such a path P from $v_{(i)}$ to $v_{(i+1)}$ exists. Suppose also that P connects to both these vertices from the same side, left or right. Consider the shifted version P' of P in S^{n+1} from $v_{(i+1)}$ to $v_{(i+2)}$. These path thus share an edge, but they must diverge at some vertex. This means there exist a vertex of degree at least 3 in S^{n+1} which is impossible since S^{n+1} is a union of two matchings. Thus P must connect to the two vertices $v_{(i)}$ to $v_{(i+1)}$ from opposite sides. This means that it crosses boundaries an even number of times. By bipartiteness the path is of even length, and together this means that the first and last edge can not be from the same matching. This implies that $v \in X_1 \cap \overline{X_1}$ or $v \in X_2 \cap \overline{X_2}$, contradicting the existence of P.

□

The following claim along with Claim 5 proves the theorem.

Claim 7. S^n does not have a blocking path for $n \geq k$.

Proof. Assume that S^n has a blocking path for $n \geq k$. This blocking path must pass each of the $n + 1$ boundaries (including left and right ends) at least once. Therefore we can find integers i and j such that this blocking path has a segment P connecting $v_{(i)}$ to $v_{(j)}$ for some $1 \leq v \leq k$. By Claim 6, we have $j > i + 1$. Now consider the graph S^{n+1}. This graph also has this path P from $v_{(i)}$ to $v_{(j)}$ and also a path P' from $v_{(i+1)}$ to $v_{(j+1)}$ that is simply a "shifted" version of P. By Claim 6, these paths are vertex disjoint. Because if they intersect then we can construct a path from $v_{(i)}$ to $v_{(i+1)}$ in S^{n+1}. By using Lemma 3 we conclude that the paths P and P' must intersect. This concludes the proof. □

□

3.3 Disjoint paths

We consider several different variants of the disjoint paths problem, but there is significant overlap in the different approaches. In each case we define a monoid $\mathcal{M}$ and show it is aperiodic. We can thus compute the word problem over $\mathcal{M}$ by AC^0 circuits and we can use these to solve the disjoint paths problem.

The monoids. We describe here the monoid in general terms. Elements of $\mathcal{M}$ consist of a (downward closed) family of sets of edges between the set of vertices $W = \{1, \ldots, w\} \cup \{1', \ldots, w'\}$. Consider a grid-planar graph G. This may be either undirected or directed. We construct a monoid-element $G^{\mathcal{M}}$ from G as follows, by letting every set of disjoint paths in G between vertices from W give rise to a set of corresponding edges in $G^{\mathcal{M}}$. Depending on the setting these paths may be vertex-disjoint or edge-disjoint, and if the graph is directed the edges are directed accordingly. The operation of the monoid will be the natural operation that makes the map $G \mapsto G^{\mathcal{M}}$ a homomorphism. Note that if $A \subseteq A'$ and $B \subseteq B'$ then $AB \subseteq A'B'$.

Reduction to monoid product. Let G be a grid-planar directed graph with pairs of terminals $(s_1, t_1), \ldots, (s_k, t_k)$. Consider the partition of G into at most $2k + 1$ segments obtained by dividing at every layer containing a terminal. For each segment we divide the graph into segments of length 1, translate these to monoid elements and compute the product of these. This results in at most $2k+1$ monoid elements describing all possible disjoint paths connecting endpoints of every segment. Since k is fixed this is a fixed amount of information from which it can then be directly decided whether disjoint paths exist between all pairs of terminals.

Showing aperiodicity of the monoid. The approach we will use in all cases is as follows. Let $\mathcal{G}$ be a group in $\mathcal{M}$ with identity E, and let A be any element of $\mathcal{G}$. We shall then prove that $E \subseteq A$. Note then that this means $A^{-1} = EA^{-1} \subseteq AA^{-1} = E$, and hence $A = E$. Showing this for all A implies that $\mathcal{G}$ is trivial.

3.3.1 Edge-disjoint paths in upward planar graphs

Here we consider *directed upward grid-planar graphs*, i.e., every edge in the graph is directed from some vertex in layer i to some vertex in layer $i + 1$ or it is directed from a layer i vertex to another layer i vertex that is at a distance of one unit on the grid from the source vertex and edges do not cross. We are to decide if edge-disjoint paths exists from $(s_1, t_1), \ldots, (s_k, t_k)$ for fixed k. Thus if we consider the monoid element corresponding to such a graph, each multiset of edges in the monoid element contains only directed edges from the left-side vertices $\{1, \ldots, w\}$ to the right-side vertices $\{1', \ldots, w'\}$, and these correspond to pairwise edge-disjoint paths in the corresponding directed upward planar grid graph. An example is shown in Figure 4.

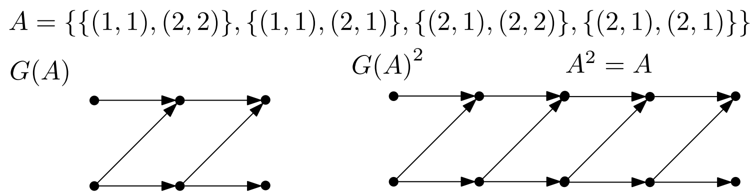

Figure 4: Monoid for Edge-disjoint Paths in Upward Planar Graphs

Our main theorem in this section is the following.

Theorem 12. *For any fixed k, given a directed upward grid-planar graph G and k pairs of vertices $(s_1, t_1), \ldots, (s_k, t_k)$, we can decide whether there are pairwise edge-disjoint paths from s_i to t_i for $i = 1$ to k in AC^0.*

The theorem follows from the following claim

Claim 8. $\mathcal{M}$ is aperiodic.

Proof. Let $\mathcal{G}$ be a group in $\mathcal{M}$ with identity E. Let A be an element of $\mathcal{G}$ such that $o_{\mathcal{G}}(A) = p \geq 2$. We are to show that $E \subseteq A$.

Let $G(E)$ and $G(A)$ be grid-planar graphs such that $G(E)^{\mathcal{M}} = E$ and $G(A)^{\mathcal{M}} = A$. Let $S \in E$, $c = |S|$ be the number of edges in S between the sets $\{1, \ldots, w\}$ and $\{1', \ldots, w'\}$, and let $t = w^c + 1$. Let these edges be $(s_1, t_1), \ldots, (s_c, t_c)$. Consider the concatenation $G(E)^{t+1}$ of the graph $G(E)$ with itself. Since $E^{t+1} = E$ we have in $G(E)^{t+1}$ disjoint paths corresponding to S, and we will think of S as those paths. Since we have t boundaries between the $t+1$ copies of $G(E)$, there must be two boundaries such that for each of the c paths, the vertices of the two boundaries that the path crosses are the same. Let $j_1 \geq j_2 \geq \cdots \geq j_c$ be these vertex numbers. Splitting the graph $G(E)^{t+1}$ at these two layers divides each path into 3 parts. Thus the middle part consists of edge-disjoint (j_i, j_i)-paths for $i = 1, \ldots, c$. The left part will contain edge-disjoint paths (s_i, j_i) and the right part will contain edge-disjoint (j_i, t_i) paths . Also note that the left part and the right part are also graphs that correspond to the monoid element E; So E must have these paths as well. We will show that $G(A)^{2pc+1}$ also contains such paths. Since $A = EA^{2pc+1}E$ we can concatenate (s_i, j_i) paths from the $G(E)$ on the left side with the concatenation of (j_i, j_i) paths in $G(A^{2pc+1})$ and (j_i, t_i) paths in the $G(E)$ on the right side to show that $S \in A$.

Since $A^p = E$ we have that $G(A)^p$ contains (j_i, j_i)-paths for $i = 1, \ldots, c$. Denote these by $P_1, \ldots, P_c$. These paths are edge-disjoint, but may cross each other at a vertex. We first argue that without loss of generality we may assume that they never cross; they may touch each other at vertices, however. Indeed, if this is not the case we can construct such paths $\widehat{P_1}, \ldots, \widehat{P_c}$ as follows. Let H be the union of $P_1, \ldots, P_c$. Suppose we have already found the paths $\widehat{P_1}, \ldots, \widehat{P_{i-1}}$. Then let $\widehat{P_i}$ be the top-most (j_i, j_i) path in H, and then erase the edges of $\widehat{P_i}$ from H.

We can think of the paths $P_1, \ldots, P_c$ as infinite paths with period p in an infinite concatenation of the graph $G(A)$. Let $P'_1, \ldots, P'_c$ be the paths obtained by shifting the paths by the length of one graph $G(A)$ to the right. By Lemma 3 we have that P_i and P'_i must intersect in $G(A)^p$. Let $Q_1, \ldots, Q_c$ be the paths such that Q_i is the upper envelope path of P_i and P'_i. We are now ready to construct the edge-disjoint (j_i, j_i) paths $R_1, \ldots, R_c$ in $G(A)^{2pc+1}$. We think of the graph $G(A)^{2pc+1}$ as $2c$ blocks of $G(A)^p$ followed by a single $G(A)$. The path R_i proceeds as follows. In the first $i-1$ blocks it follows P_i. Then in block i, when P_i intersects Q_i it follows Q_i, and continues to do so for the following $2(c-i)$ blocks. Then in block $2c-i+1$ when P'_i intersects Q_i it follows P'_i, and continues to do so for the remaining $i-1$ blocks. After these $2c$ blocks, the c paths that started out as $P_1, \ldots, P_c$ are ending as $P'_1, \ldots, P'_c$ and they follow $P'_1, \ldots, P'_c$ through the last graph $G(A)$, thereby making Q_i a (j_i, j_i)-path. We claim that the paths $R_1, \ldots, R_c$ are edge-disjoint. Observe the path R_i. In the first $i-1$ blocks it cannot intersect any of the paths $R_1, \ldots, R_{i-1}$. This is because $R_i = P_i$ for these blocks and R_j for $j < i$ is either P_j in which case R_i is disjoint from these as P_i is disjoint from P_j for all j or R_j is the upper envelope of P_j and P'_j in which case R_j can only move further away from R_i. After that all R_is start following Q_is and they remain edge-disjoint as Q_is are edge-disjoint. Now note that R_i switches to P'_i before R_j for $j < i$. So R_i cannot intersect with R_j as R_i can only move further away from R_j (which is still following the upper envelope). Now once all R_is have switched to P'_is they remain edge-disjoint as P'_is are edge-disjoint.

□

3.3.2 Vertex-disjoint paths

Our main theorem in this section is the following.

Theorem 13. *For any fixed k, given a directed grid-planar graph G and k pairs of vertices $(s_1, t_1), \ldots, (s_k, t_k)$, we can decide whether there are pairwise vertex-disjoint paths from s_i to t_i for $i = 1$ to k in* AC^0.

A monoid element $G^{\mathcal{M}}$ consists of a family of sets of directed edges on the set $\{1, \ldots, w\} \cup \{1', \ldots, w'\}$ that form partial matchings, that is no two edges share an endpoint. Note that the edges could go from a vertex to a vertex on the same side (For ex., from 1 to 2 or from $1'$ to $2'$).

The theorem follows from the following claim.

Claim 9. $\mathcal{M}$ is aperiodic.

Proof. Let $\mathcal{G}$ be a group in $\mathcal{M}$ with identity E. Let A be an element of $\mathcal{G}$ such that $o_{\mathcal{G}}(A) = p \geq 2$. We are going to show that $E \subseteq A$.

Let $G(E)$ and $G(A)$ be grid-planar graphs such that $G(E)^{\mathcal{M}} = E$ and $G(A)^{\mathcal{M}} = A$. Let $S \in E$ and let $t = w^{w+1} + 1$. We shall prove that $S \in A$. Let c be the number of edges in S between the sets $\{1, \ldots, w\}$ and $\{1', \ldots, w'\}$. We call the corresponding paths crossing paths.

Consider the concatenation $G(E)^{t+1}$ of the graph $G(E)$ with itself. Since $E^{t+1} = E$ we have in $G(E)^{t+1}$ disjoint paths corresponding to S. We will think of S also as this set of paths. This set will induce vertex disjoint paths in each of the $t+1$ copies of $G(E)$. We will now have two cases: Either all the $t+1$ graphs have exactly c crossing paths or some graph has $c' > c$ crossing paths. In case some graph $G(E)$ has $c' > c$ crossing paths, let S' be the set of paths induced by S in that graph. We then start over, and prove that $S' \in A$. This will imply that $S \in A$ since $A = EAE$, and this case can only occur a finite number of times since $c' \leq w$.

So we may now suppose that all graphs have exactly c crossing paths. This means also that each path in S crosses each graph $G(E)$ of the concatenation $G(E)^{t+1}$ using a single crossing path. Between two such crossing paths, a path (in S) may cross the boundary between two graphs a number of times. Record for each boundary and for each path an ordered list of the vertex numbers in which the path crosses the boundary. By the choice of t there must be two boundaries for which each path cross in the same way in the two boundaries. We then consider the part of the concatenation between these two boundaries and let S' be the disjoint paths induced on this part. Clearly $S' \in E$ as we have taken S' from a concatenation of $G(E)$ graphs, and we will prove that $S' \in A$ and be done again since $A = EAE$.

We have a set of disjoint paths corresponding to S' in the graph $G(A)^p$ (Since $A^p = E$). We can think of these paths as being induced by infinite paths $P_1, \ldots, P_c$ with period p in an infinite concatenation of the graph $G(A)$. We suppose these are ordered such that P_1 is the top-most path, P_2 the second top-most path, and so on. Let $P'_1, \ldots, P'_c$ be the infinite paths obtained by shifting the paths by the length of one graph $G(A)$ to the right. By Lemma 3 we have that P_i and P'_i must intersect in $G(A)^p$. Let $Q_1, \ldots, Q_c$ be the paths such that Q_i is the upper envelope path of P_i and P'_i.

We are now ready to construct vertex-disjoint paths $R_1, \ldots, R_c$ corresponding to S in $G(A)^{(2c+2)p+1}$. As before we consider $G(A)^{(2c+2)p+1}$ as $2c+2$ blocks of $G(A)^p$ followed by a single $G(A)$. Path R_i proceeds as follows. In the first block it follows exactly along P_i, and it continues to do so for the next $i-1$ blocks. Then in block $i+1$, when P_i intersects Q_i it follows Q_i. Note that this may lead R_i into the previous block, but it will not intersect itself. From block $i+1$ the path R_i continues to follow along Q_i for the following $2(c-i)$ blocks. Then in block $2c-i+2$ when P'_i intersects Q_i it follows P'_i and continues to do so for the next $i-1$ blocks. After these $2c+1$ blocks the paths R_i continue to follow along P'_i for another block, and also through the last graph $G(A)$. Using an argument similar to the one we used in proving Claim 8, we can conclude that $R_1, \ldots, R_c$ are pairwise vertex-disjoint and this completes the proof. □

3.3.3 Edge-disjoint paths in undirected graphs

Here we consider the setting where the graphs are undirected, and we are to decide if edge-disjoint paths exists. Thus the monoid elements consists of sets of undirected edges. The proof will use ideas from both of the two previous paragraphs.

Aperiodicity of the monoid. Let $\mathcal{G}$ be a group in $\mathcal{M}$ with identity E. Let A be an element of $\mathcal{G}$ of period p. We are to show that $E \subseteq A$.

The proof first proceeds exactly as in the vertex-disjoint case. Thus we let $G(E)$ and $G(A)$ be grid-planar graphs such that $G(E)^{\mathcal{M}} = E$ and $G(A)^{\mathcal{M}} = A$. Let $S \subseteq E$ and let c be the number of cross-edges in S. By considering the concatenation $G(E)^{t+1}$ with $t = w^{w+1} + 1$, we may reduce to the case where we have a set of disjoint paths corresponding to S in the graph $G(A)^p$, and where we can this of these paths as being induced by infinite paths $P_1, \ldots, P_c$ with period p in an infinite concatenation of the graph $G(A)$.

Now, we depart from the similarity with the vertex-disjoint case. The infinite paths $P_1, \ldots, P_c$ are just assumed to be edge-disjoint so they may intersect at vertices. From these we shall now construct infinite paths $\widehat{P_1}, \ldots, \widehat{P_c}$ that are edge disjoint, but may only touch at vertices, never cross. From the construction the paths are ordered from top path $\widehat{P_1}$ to the bottom path $\widehat{P_c}$. Let H be the union of all the paths infinite paths $P_1, \ldots, P_c$. Suppose we have already found the paths $\widehat{P_1}, \ldots, \widehat{P_{i-1}}$. Let $\widehat{P_i}$ be the upper-envelope path in H, and then erase the edges of $\widehat{P_i}$ from H. (This is the step that does not generalize to the case directed graphs). Assume from now on that the original paths $P_1, \ldots, P_c$ are ordered in such a way that P_i intersects $\widehat{P_i}$.

Let $P'_1, \ldots, P'_c$ be the infinite paths obtained by shifting the paths $P_1, \ldots, P_c$ by the length of one graph $G(A)$ to the right. Similarly, let $\widehat{P_1}', \ldots, \widehat{P_c}'$ be the infinite paths obtained by shifting the paths $\widehat{P_1}, \ldots, \widehat{P_c}$ by the length of one graph $G(A)$ to the right. Let $Q_1, \ldots, Q_c$ be the paths such that Q_i is the upper envelope path of $\widehat{P_i}$ and $\widehat{P_i}'$.

We are now ready to construct new paths $R_1, \ldots, R_c$. These are constructed in blocks of $G(A)^p$, and follows several phases. We start out with the paths $P_1, \ldots, P_c$. These are followed for one block. Then we have a transition phase to the paths $\widehat{P_1}, \ldots, \widehat{P_c}$. This gives the geometric ordering and in the next two transition phases we transition to the shifted versions $\widehat{P_1}', \ldots, \widehat{P_c}'$ via the upper envelope paths $Q_1, \ldots, Q_c$. Finally in the last transition phase we transition back to the original (but shifted) paths $P'_1, \ldots, P'_c$. These are then followed through one block and then one more graph $G(A)$ to complete the paths.

We shall describe the transition between the paths $P_1, \ldots, P_c$ and $\widehat{P_1}, \ldots, \widehat{P_c}$, and the rest of the proof then follows analogously to the previous settings. Here we are in the situation that the paths $R_1, \ldots, R_c$ have followed the paths $P_1, \ldots, P_c$ for one block. Then in the next block, when P_1 first intersects with $\widehat{P_1}$ we will let R_1 follow along $\widehat{P_1}$, and we do this by also modifying the other paths $R_2, \ldots, R_c$ accordingly. More precisely, whenever two or more paths meet at a later vertex of $\widehat{P_1}$ we exchange the continuations of curves if necessary in order to let R_1 follow along $\widehat{P_1}$. We continue in a similar way in the next $c-1$ blocks, letting R_j start following along $\widehat{P_j}$ beginning from block j.

Performing all the transitions as described above we have shown that $S \in A^{p(1+4c)+1} = A$, thereby completing the proof.

4 Lower bounds

In this section we show hardness for the problems studied in the previous section. All the lower bounds are under projection reductions. We will thus given a circuit C build a graph G with edges labeled by literals, i.e. variables or negations of variables. Given an input x, let $G(x)$ be the graph obtained from G by keeping exactly the edges labeled by literals that are 1 under the assignment x. We then show that $C(x) = 1$ if and only if the graph $G(x)$ satisfies the graph property under consideration.

Our NC^1 lower bounds for the non-planar case build the characterization of NC^1 in terms of permutation branching programs by Barrington[2]. His construction gives for any polynomial size NC^1 circuit a polynomial length *program* over the group S_5 of permutations of 5 elements. From a program of length l we can construct a graph with vertices placed on the grid $\{1, \dots, l+1\} \times \{1, \dots, 5\}$. Between two layers of the grid we have 10 edges corresponding to the two permutations corresponding to an instruction of the program. These are then labeled by the corresponding variable or its negation accordingly. The resulting graph $G(x)$ will for any input consist of exactly 5 disjoint paths, and we can without loss of generality assume these are as shown in Figure 5 (a) and (b). We can without loss of generality assume that the length $l + 1$ is even or odd if needed. We shall denote this graph the Barrington graph.

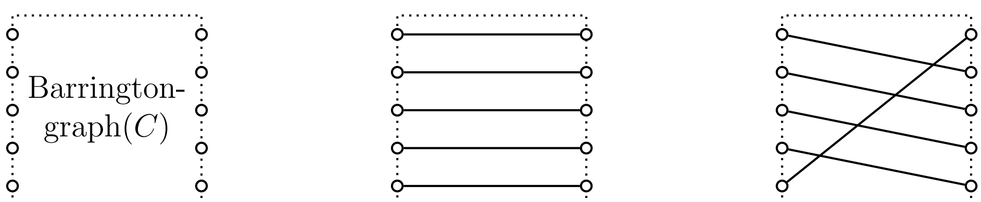

(a) (b)

Figure 5: Graph for a NC^1 circuit C and paths when $C(x) = 0$ (a) and $C(x) = 1$ (b).

Barrington et al.[3] showed that connectivity in constant width grid graphs is complete for AC^0 under projection reductions. This holds for both undirected graphs and directed graphs. Thus for any AC^0 circuit C we have a grid graph G with edges labeled by literals, with a vertex s in the first layer, a vertex t in the last layer, such that $G(x)$ has a (s, t) path if and only if $C(x) = 1$. We can also here without loss of generality assume the length of the grid graph is even or odd if needed. We shall denote this graph the BLMS graph.

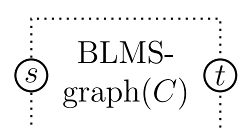

4.1 2-coloring

The results here are based upon the simple fact that a cycle can be 2-colored if and only if it is of even length, and that if a path is 2-colored then the endpoints must have different color if and only

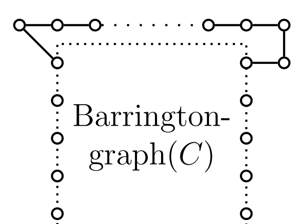

if the path is of odd length. In the non-planar case we obtain NC^1-hardness simply by taking the Barrington graph and connecting the top nodes on the two sides by a path whose length is of opposite parity of the length of the graph. Then if these nodes are connected in the Barrington graph an odd cycle appears, and this happens exactly when $C(x) = 0$. Otherwise the graph consists of disjoint paths and can thus be 2-colored.

We next consider the planar case. From the BLMS-graph for a given AC^0 circuit C we construct a 2-coloring gadget graph. The graph is always 2-colorable (since the BLMS graph is bipartite) and has even length. Also if $C(x) = 1$ then s and t must have different colors in any 2-coloring, since in that case there is a path of odd length (through the BLMS graph for C) between s and t. Similarly, if $C(x) = 0$, then s and t must have the same color in any 2-coloring, since in that case there is a path of even length (through the BLMS graph for $\neg C$) between s and t. Consider next a $\mathsf{XOR} \circ \mathsf{AC}^0$ circuit C. Suppose that $C = \mathrm{XOR}(C_1, \ldots, C_m)$ where $C_1,\ldots,C_m$ are AC^0 circuits. We simply concatenate the 2-coloring gadgets for $C_1, \ldots, C_m$ as shown in figure 6 and connect the s terminal of the first gadget graph with the t terminal of the last gadget graph by an odd length path.

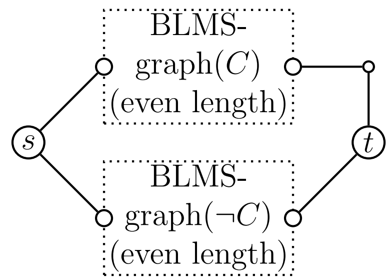

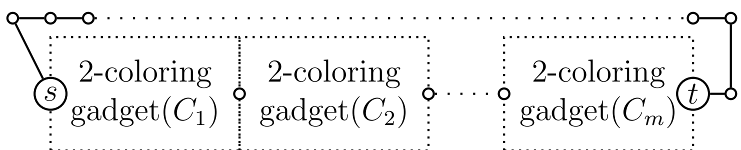

Figure 6: Graph for $\mathrm{XOR}(C_1, \ldots, C_m)$, where $C_1,\ldots,C_m$ are AC^0 circuits.

By the property of the gadget graphs, in any 2-coloring, the vertices s and t must have different color exactly when an odd number of the circuits $C_1, \ldots, C_m$ evaluate to 1, and must have the same color otherwise. Since the top path connecting s and t has odd length it follows that the graph can be 2-colored exactly when $\mathrm{XOR}(C_1(x), \ldots, C_m(x)) = 1$. Now given an $\mathsf{AND} \circ \mathsf{XOR} \circ \mathsf{AC}^0$ circuit $C = \mathrm{AND}(C_1, \ldots, C_m)$ we construct the above graph for each of the $\mathsf{XOR} \circ \mathsf{AC}^0$ sub-circuits consider the graph that is the disjoint union of all these graphs. Then this graph can be 2-colored if and only if $C(x) = 1$.

4.2 3-coloring

To show NC^1-hardness in the non-planar case we can adapt the graph constructed for the case of 2-coloring, and change it appropriately. Namely we just replace each edge, except for the rightmost vertical edge, by the simple 4-vertex gadget graph 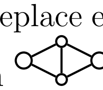. The gadget graph ensures that all of the original vertices must be of the same color, and hence a coloring is not possible because of the rightmost vertical edge if a cycle is formed in the original graph.

We can transform this graph into a grid graph with diagonals by using the simple crossover gadget as shown on the right that was constructed by Garey et al.[11] for showing that the 3-coloring problem for general planar graphs is NP-complete. This gadget simulates by a grid graph with diagonals a crossing between edges (a, b) and (c, c). Since the graph we start with is layered we can deal with the intersections in each layer separately. By appropriately placing a number of crossover gadgets we can simulate the crossings between the layers by connecting vertices of the surrounding layers to the crossing gadgets. To make the entire graph be a grid graph with diagonals we can finally replace such edges by concatenations of the simple 4-vertex gadget graph together with a regular edge. Doing this shows that 3-coloring remains NC^1-hard for constant width grid graphs with diagonals.

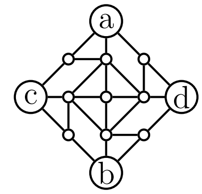

4.3 Matching

The results here are based upon the simple fact that a path has a perfect matching if and only if it is of odd length. In the non-planar case we obtain NC^1-hardness simply by taking the Barrington graph of odd length and attaching additional edges to the top vertex on the left side and the second vertex from the top on the right hand side. Then the graph has a perfect matching if and only if these two vertices are connected by a path through the Barrington graph, which in turn happens if and only if $C(x) = 1$. We next consider the planar case. As in the case of 2-coloring we shall for an AC^0 circuit C build a gadget based on the BLMS graph. But we shall first make a modification to the BLMS graph to make it suited for the purpose of a gadget for matching.

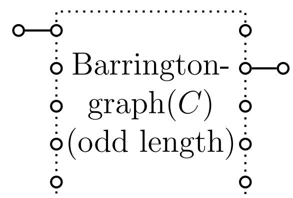

We start with the directed version of the BLMS graph. We shall also assume it to be of even length. From this we shall construct an undirected graph we denote by BLMS*. It is obtained by performing the standard reduction from directed (s, t) connectivity to perfect matching. Namely, we split each vertex u into two vertices u and u', connected by an edge. Then for every directed edge from u to v we connect u' and v. We remove the vertices s and t'. The graph always has a partial matching where only s' and t are left unmatched. Furthermore, the graph has a perfect matching if and only if $C(x) = 1$. We can do this preserving planarity and such that the resulting graph has even length. Note also that is a bipartite graph, and we therefore obtain AC^0-hardness for perfect matching on bipartite grid-planar graphs. With some more work one can also obtain an equivalent grid graph.

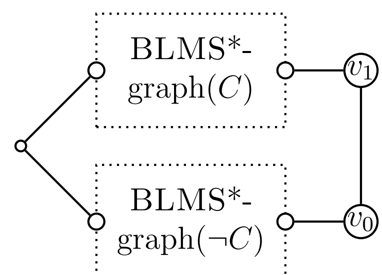

From the BLMS* graphs we construct a matching gadget graph for the circuit C. Let us call all vertices except for v_1 and v_0 for internal vertices. In this graph there is always a partial matching that matches all internal vertices and exactly one of v_0 or v_1. Also if $C(x) = 1$ then any partial matching that matches all internal vertices must leave v_1 unmatched. Similarly, if $C(x) = 0$ then any partial matching that matches all internal vertices must leave v_0 unmatched. We construct the matching gadget graph in such a way it is of even length. Consider next a $\mathsf{XOR} \circ \mathsf{AC}^0$ circuit C. Suppose that $C = \text{XOR}(C_1, \ldots, C_m)$ where $C_1,\ldots,C_m$ are AC^0 circuits and assume without loss of generality that m is even. We place the matching gadgets for $C_1, \ldots, C_m$ adjacent to each other as shown in Figure 7. We also have a top path ending in a terminal v_1 and a bottom path ending in a terminal v_0 along the gadgets. The top path is of constructed to be of odd length and the bottom path of even length. The v_1 terminal of a gadget graph is connected to the top path and the v_0 is connected to the bottom path. They may thus "steal" a vertex from either the top or bottom path depending on the unmatched terminal. In other words, in order to match all the vertices of a gadget for C_i, in case $C_i(x) = 1$ the gadget must steal a vertex from the top path, and in case $C_i(x) = 0$ the gadget must steal a vertex from the bottom path. We can see that the combined graph always have a partial matching where all vertices except exactly one terminal is matched. If $C(x) = 1$ an odd number of vertices are stolen from both the top and bottom path. Thus if all vertices besides the terminals are matched then v_1 is unmatched and v_0 is matched. Similarly, if $C(x) = 0$ an even number of vertices are stolen from both the top and bottom paths. Thus if all vertices besides the terminals are matched then v_1 is matched and v_0 is unmatched. Thus for C we get a matching gadget similar to the matching gadget constructed for AC^0 circuits.

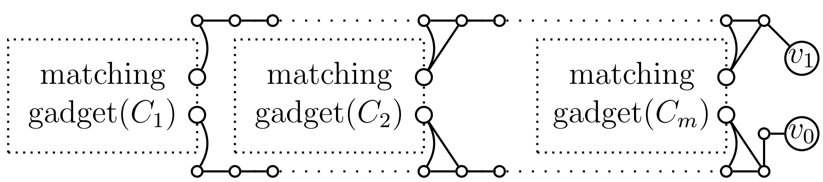

Figure 7: Graph for $\text{XOR}(C_1, \ldots, C_m)$, where $C_1,\ldots,C_m$ are AC^0 circuits (even m).

As opposed to the case of 2-coloring we can here go a step further. Consider now an $\mathsf{OR}{\circ}\mathsf{XOR}{\circ}\mathsf{AC}^0$ circuit C. Suppose that $C = \text{XOR}(C_1, \ldots, C_m)$ where $C_1,\ldots,C_m$ are $\mathsf{XOR}{\circ}\mathsf{AC}^0$ circuits and assume without loss of generality that m is even. We place the matching gadgets for $C_1, \ldots, C_m$ adjacent to each other as shown in Figure 8. As before we have a top path and a bottom path along the gadgets. Both of the paths are of even length. Different to before, the gadgets are constructed in such a way that if the terminal v_1 is unmatched the gadget may steal a vertex from *either* the top path or the bottom path. If the terminal v_0 is unmatched the gadget may just steal a vertex from the bottom path. Since both the top and bottom path are of even length, the only way to match all vertices is that the gadgets steal an odd number of vertices from both paths. Now if $C(x) = 1$ at least one gadget can be matched such that v_1 is left unmatched. Thus we may pick an odd number of subcircuits C_i to steal a vertex from the top path, and let the remaining steal a vertex from the bottom path. On the other hand if all $C_i(x) = 0$ then to match all vertices of each gadget they need to steal a vertex from the bottom path, meaning that the full graph does not have a perfect matching.

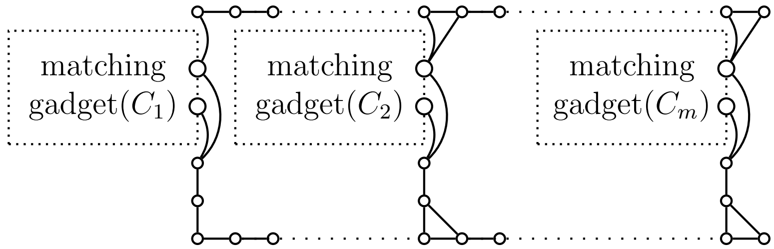

Figure 8: Graph for OR($C_1, \ldots, C_m$), where $C_1,\ldots,C_m$ are $\mathsf{XOR} \circ \mathsf{AC}^0$ circuits (even m).

As in the case of 2-coloring, if we construct the final graph to be a disjoint union of such graphs, the matching problem on the resulting graph simulates $\mathsf{AND} \circ \mathsf{OR} \circ \mathsf{XOR} \circ \mathsf{AC}^0$ circuits.

4.4 Hamiltonian cycle

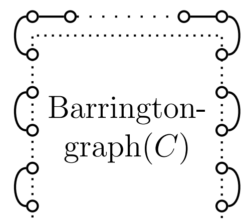

In the non-planar case we obtain NC^1-hardness simply by taking the Barrington graph, add another path on top, and connecting the nodes on the two sides in pairs. If $C(x) = 0$ the resulting graph consists of 3 disjoint cycles, whereas if $C(x) = 1$ the resulting graph consists of a single cycle. We can not similarly to the case of 3-coloring translate this directly to the planar case, since we have no general crossover gadget. We can however proceed following the approach shown for establishing NP-completeness for the case of general planar graphs. We shall just outline the proof. Plesńik[15] showed that the Hamiltonian Cycle problem over directed planar graphs of degree 2 is NP-complete. Our main observation is that we can use the same reduction to reduce the complement of directed connectivity the Hamiltonian Cycle problem on bounded width planar graphs. Given a graph G with source s and target t we derive a 2-CNF formula F consisting of the following clauses: (s), $(\neg t)$ and for each directed edge $(u \to v)$ there is a clause $(\neg u \vee v)$. It follows that F is satisfiable if and only if there is no path in G from s to t. We then apply the reduction in [15] to F. We observe that the resulting graph is planar and bounded width. To see this, we note that clauses in F can be ordered so that for every node u, the set of clauses containing u appear in an interval of constant length. This order essentially follows the order according to G. After applying the reduction to F at some point a graph with crossings is obtained. But because of the interval property of F all crossings can appear in such intervals of constant length, which means that we can apply the crossing gadget to each such interval and blow up the width by only a constant. To get a grid graph we then apply a reduction from [13]. We first transform the graph obtained in the first part of the reduction to a bipartite graph and then embed it into a grid. However we should note that [13] embeds a n-vertex graph into a $\Theta(n) \times \Theta(n)$ grid. Since the graph that we start with is layered, we can apply the embedding on consecutive layers separately and hence we will get a constant width grid graph.

4.5 Disjoint paths

Here we need just remark that the disjoint paths problem is precisely the connectivity problem when $k = 1$. Thus we have NC^1-hardness in the non-planar case by the result of Barrington[2] and AC^0-hardness in the planar case by the result of Barrington et al.[3].

5 Discussion & Open Problems

In this paper, we studied that the circuit complexity of several computational problems on graphs of constant planar cutwidth. We used the Barrington-Therein characterization of fine structure of NC^1 to establish the circuit complexity upper bounds. The recipe for the the proof was simply as follows, define an appropriate monoid from the setting of the problem and establish that the solution of the computational problem is exactly equivalent to solving the word problem over this monoid. As the next step, we analyse the algebraic structure of the monoid and establish various properties.

The main open problems that arises from our paper is closing the gap in the complexity of perfect matching and 2-coloring problems as mentioned in table in figure 1. The applicability of the framework to determine the circuit complexity of more computational problems is also another potential direction of study.

References

[1] Eric Allender and Ulrich Hertrampf. Depth reduction for circuits of unbounded fan-in. *Information and Computation*, 112(2):217–238, 1994.

[2] David A. Mix Barrington. Bounded-width polynomial-size branching programs recognize exactly those languages in NC^1. *J. Comput. Syst. Sci*, 38(1):150–164, 1989.

[3] David A. Mix Barrington, Chi-Jen Lu, Peter Bro Miltersen, and Sven Skyum. Searching constant width mazes captures the AC^0 hierarchy. In *STACS*, volume 1373 of *Lecture Notes in Computer Science*, pages 73–83. Springer, 1998.

[4] David A. Mix Barrington and Denis Thérien. Finite monoids and the fine structure of NC^1. *J. ACM*, 35(4):941–952, 1988.

[5] H. L. Bodlaender. Some classes of planar graphs with bounded treewidth. *Bulletin of the EATCS*, 36:116–126, 1988.

[6] Marek Cygan, Dániel Marx, Marcin Pilipczuk, and Michal Pilipczuk. The planar directed k-vertex-disjoint paths problem is fixed-parameter tractable. In *FOCS*, pages 197–206. IEEE Computer Society, 2013.

[7] Samir Datta, Arjun Gopalan, Raghav Kulkarni, and Raghunath Tewari. Improved bounds for bipartite matching on surfaces. In *STACS*, volume 14 of *LIPIcs*, pages 254–265. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2012.

[8] Michael Elberfeld, Andreas Jakoby, and Till Tantau. Algorithmic meta theorems for circuit classes of constant and logarithmic depth. In *STACS*, volume 14 of *LIPIcs*, pages 66–77. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2012.

[9] Walter Feit and John G. Thompson. Solvability of groups of odd order. *Pacific J. Math.*, 13(3):775–1029, 1963.

[10] Steven Fortune, John E. Hopcroft, and James Wyllie. The directed subgraph homeomorphism problem. *Theor. Comput. Sci*, 10:111–121, 1980.

[11] M. R. Garey, David S. Johnson, and Larry J. Stockmeyer. Some simplified NP-complete graph problems. *Theor. Comput. Sci*, 1(3):237–267, 1976.

[12] M. R. Garey, David S. Johnson, and Robert Endre Tarjan. The planar hamiltonian circuit problem is NP-complete. *SIAM J. Comput*, 5(4):704–714, 1976.

[13] Alon Itai, Christos H. Papadimitriou, and Jayme Luiz Szwarcfiter. Hamilton paths in grid graphs. *SIAM J. Comput*, 11(4):676–686, 1982.

[14] P. W. Kasteleyn. Graph theory and crystal physics. In F. Harary, editor, *Graph Theory and Theoretical Physics*, pages 43–110. Academic Press, 1967.

[15] Ján Plesník. The NP-completeness of the hamiltonian cycle problem in planar digraphs with degree bound two. *Inf. Process. Lett*, 8(4):199–201, 1979.

[16] Alexander A. Razborov. Lower bounds for the size of circuits of bounded depth with basis $(\wedge, \oplus)$. *Mathematical Notes of the Academy of Science of the USSR*, 41(4):333–338, 1987.

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

[18] Neil Robertson and Paul D. Seymour. Graph minors.XIII. the disjoint paths problem. *J. Comb. Theory, Ser. B*, 63(1):65–110, 1995.

[19] Alexander Schrijver. Finding k disjoint paths in a directed planar graph. *SIAM J. Comput*, 23(4):780–788, 1994.

[20] Vijay V. Vazirani. NC algorithms for computing the number of perfect matchings in $K_{3,3}$-free graphs and related problems. *Inf. Comput*, 80(2):152–164, 1989.