Planarity, Determinants, Permanents, and (Unique) Matchings

Samir Datta[1], Raghav Kulkarni[⋆2], Nutan Limaye[3], and Meena Mahajan[3]

[1] Chennai Mathematical Institute, Siruseri, Chennai 603 103, India.
sdatta@cmi.ac.in
[2] Dept. of Computer Science, Univ. of Chicago, U.S.A. raghav@cs.uchicago.edu
[3] The Institute of Mathematical Sciences, Chennai 600 113, India.
{nutan,meena}@imsc.res.in

Abstract. We explore the restrictiveness of planarity on the complexity of computing the determinant and the permanent, and show that both problems remain as hard as in the general case, i.e. GapL and #P complete. On the other hand, both bipartite planarity and bimodal planarity bring the complexity of permanents down (but no further) to that of determinants. The permanent or the determinant modulo 2 is complete for ⊕L, and we show that parity of paths in a layered grid graph (which is bimodal planar) is also complete for this class. We also relate the complexity of grid graph reachability to that of testing existence/uniqueness of a perfect matching in a planar bipartite graph.

1 Introduction

For many natural problems on graphs, imposing planarity does not reduce the complexity. For instance, vertex cover is NP-complete, and remains so even for planar degree-3 restrictions; so does planar 3-dimensional matching [15]. The circuit value problem is P-complete, and remains so even if the graph underlying the circuit is restricted to be planar. In [19] and [27], the complexity of several counting problems has been investigated under planar restrictions. More recently, [32] establishes that counting vertex covers remains #P-complete even when restricted to 3-regular planar bipartite graphs. Thus there is some evidence to believe that planarity is not a real restriction at all.

However, there are notable exceptions. In the circuit setting, for instance, monotone circuit value is P-complete, but monotone planar circuit value is in NC [33, 14]. Constant-width circuits characterize NC^1 [7], while planar constant-width circuits characterize its subclass ACC^0 [16]. In the graph-theoretic setting, counting the number of perfect matchings in a bipartite graph is #P-hard [28], while counting it in a planar bipartite graph (or even in a planar non-bipartite graph) is in NC [30, 21]. Another very recent exception has to do with reachability. Given a directed graph G and two vertices s and t, determining whether

⋆ Part of this work was done while this author was visiting the Chennai Mathematical Institute.

there is a path from s to t is the canonical complete problem for nondeterministic logspace NL. However, if the graph is planar, then a recent result from [9], building on the techniques of [25, 4], shows that the presence, and even the absence, of a path can be detected in unambiguous logspace UL. While UL is known to coincide with NL in the non-uniform setting, and even in the uniform setting under a plausible hardness condition [6], as of now they are not known to coincide unconditionally. So the result of [9] is an instance of planarity reducing the complexity of a problem.

Thus we see that the condition of planarity could be exploited in establishing better upper bounds in some cases. Motivated by the need to better understand how planarity can help, we examine the complexity of determinant, permanent, and unique perfect matchings when restricted to planar instances. Recall that both the determinant and the permanent of the adjacency matrix of a graph G count the total weight of all cycle covers in G, with the one difference that the determinant considers the *signed* weight. Computing the determinant (over integers or rationals) is known to be GapL-complete [13, 26, 29, 31], while computing the permanent is known to be #P-complete (see [28]; the 0-1 permanent equals the number of perfect matchings in a related bipartite graph). However, testing whether the 0-1 permanent is zero is in P and thus significantly easier than #P, whereas testing whether the 0-1 determinant is zero is complete for the exact-counting-in-logspace class $\mathsf{C_{=}L}$ [3], and thus at least as hard for NL. Interestingly, the permanent mod 2 equals the determinant mod 2 and is thus easy to compute, in fact complete for the parity logspace class $\oplus$L. Another complete problem for $\oplus$L is checking whether the number of $s \rightsquigarrow t$ paths in a directed acyclic graph is odd. Testing whether a bipartite graph has a perfect matching, B-PM, is known to be hard for NL [11], while testing whether a bipartite graph has a unique perfect matching, B-UPM, is known to be hard for NL and in $\mathsf{C_{=}L} \cap \mathsf{NL}^{\oplus \mathsf{L}}$ [18].

We examine planar restrictions of these and related problems. Our main results are summarized in Table 1. (The involved terms are explained in the respective sections.)

This paper is organised as follows. Section 2 describes the notation needed to describe the results of the paper. Sections 3 and 4 describe the hardness and the membership results respectively concerning determinant and permanent. Section 5 describes the hardness of $\oplus$LGGR for $\oplus$L, and Section 6 describes the results concerning planar B-PM and B-UPM.

2 Notation and Preliminaries

L and P denote deterministic logspace and polynomial time computation, respectively. We consider the nondeterministic classes NP and NL, their counting counterparts #P and #L, and the closures of these under subtraction GapP and GapL. We also consider (1) the exact counting in logspace class $\mathsf{C_{=}L}$; a language L is in $\mathsf{C_{=}L}$ if and only if some GapL function vanishes exactly on strings in L, and (2) the parity logspace class $\oplus$L; L is in $\oplus$L if and only if some GapL

Problem	General bound	Restriction	Our New Bound
Total signed weight of cycle covers (Determinant of adjacency matrix)	GapL-complete	planar	GapL-hard
Total weight of cycle covers (Permanent of adjacency matrix)	#P-complete	planar	#P-hard
		planar bipartite	GapL-complete
		planar bimodal	GapL-complete
Total weight of perfect matchings (Permanent of bip-adjacency matrix)	#P-complete	planar bipartite	GapL-complete
Parity of $\#s \rightsquigarrow t$ paths in directed acyclic graph	$\oplus$L-complete	planar, even layered grid graph	$\oplus$L-hard
Bipartite UPM	NL-hard, in $\mathsf{C_{=}L} \cap \mathsf{NL}^{\oplus \mathsf{L}}$	planar	in $\oplus$L, L-hard, co-LGGR-hard, equiv to GGUPM
Bipartite PM	NL-hard	planar	L-hard, GGR-hard, equiv to GGPM

Table 1. Main results

function takes odd values exactly on strings in L. It is known that $\mathsf{NL} \subseteq \mathsf{C_{=}L}$ and that $\oplus\mathsf{L}^{\oplus\mathsf{L}} = \oplus\mathsf{L}$. The canonical complete problem for NL is Reachability in a directed acyclic graph. A complete problem for GapL is computing the determinant of an integer matrix; hence testing singularity of a matrix is complete for $\mathsf{C_{=}L}$. See for instance [1].

We consider planar graphs specified by planar combinatorial embeddings: such an embedding specifies, for each vertex, the cyclic ordering of edges incident on it in some plane drawing. Testing planarity and obtaining planar combinatorial embeddings can be done in L by the results of [5, 24]. A planar embedding of a directed graph is bimodal if at every vertex, all the incoming edges appear contiguously in the cyclic ordering. Not every planar graph has a bimodal embedding. See for instance [23].

A grid graph is a directed graph with vertices laid out on the plane at integer coordinates, and edges going unit distance east-west or north-south only. A grid graph is layered if all horizontal edges are in the same direction (say left-to-right, or x-monotone), and so are all vertical edges (y-monotone).

We will frequently use the following observation:

Proposition 1. *A bipartite graph can be drawn on the plane with straight-line edges, and with no two crossings sharing the same coordinates. The combinatorial embedding corresponding to such a drawing can be obtained in logspace.*

(To draw $K_{n,n}$, place vertices of the first part on the x-axis, vertex u_i at $(0, i)$. Place vertices of the second part on the $x = 1$ line suitably spaced apart; place vertex v_j at $(1, n^{2j})$.)

For any directed graph H with a special source vertex s and sink vertex t, define the split graph Split(H) as follows: (1) split every node v into two nodes, v_{in} and v_{out}, (2) for every edge (u, v) in the original graph, draw an edge from

u_{out} to v_{in}, with the same weight, (3) draw the edges from v_{in} to v_{out} for each v, with weight 1, and (4) delete s_{in} and t_{out}; rename s_{out} and t_{in} as s and t. Note that Spilt(H) is always bipartite. Further, if H has a bimodal planar embedding, then Split(H) is also bimodal planar, and the witnessing embedding can be easily obtained from that of H. (If H is planar but not bimodal, then Split(H) may not be planar at all.)

Corresponding to any $n \times n$ matrix M, we can associate two graphs: G_M is a directed graph on n vertices, with edge $\langle i, j \rangle$ having weight $M(i, j)$, and H_M is an undirected bipartite graph on $2n$ vertices, with edge $(i, n + j)$ having weight $M(i, j)$. M is said to be the adjacency matrix of G_M and the bipartite adjacency matrix of H_M. A cycle cover in a graph is a collection of vertex disjoint cycles spanning the graph. The determinant of a matrix M, $\mathsf{Det}(M)$, equals the total signed weight of all cycle covers in G_M, while its permanent, $\mathsf{Perm}(M)$, equals the total unsigned weight of all cycle covers in G_M. The sign of a cycle cover is $(-1)^k$, where k is the number of even length cycles in the cover. $\mathsf{Perm}(M)$ also equals the total weight of all perfect matchings in H_M. Here the weight of a cycle cover or matching is the product of the weights of its constitutent edges.

3 Planarizing the Determinant and the Permanent: retaining hardness

Computing the determinant (over integers) is known to be GapL-complete [13, 26, 29, 31]. We show that it remains hard if the matrix is restricted to be the adjacency matrix of a planar graph. Weights in {0,1} suffice, and if the graph is required to be bipartite then weights in {-1,0,1} suffice. Further, a natural complete problem for GapL is DAG-WT-$s \rightsquigarrow t$-PATHS: finding the total weight of all $s \rightsquigarrow t$ paths in a weighted directed acyclic graph DAG. We show that this problem remains GapL-hard even restricting the DAG to be planar, if we allow negative weights.

We also investigate the complexity of the planar permanent. The permanent itself is #P-complete, though the hardness is under Turing reductions. There are two types of planar restrictions we can consider, and they have quite a different flavour. We want to compute $\mathsf{Perm}(M)$ when either the graph G_M or the graph H_M (see Section 2) is planar. If we require H_M to be planar, then #P-hardness is lost, because the total weight of perfect matchings in a planar (bipartite or otherwise) graph can be done in GapL using the framework of Pfaffians; see [30, 21]. We show that this is in fact not just in GapL but also GapL-complete. Though [21] shows that computing the Pfaffian is GapL-complete, the underlying graphs are not planar. We show hardness without recourse to Pfaffians.

If we require that the graph G_M is planar, then we are counting the total weight of cycle covers in a planar graph. We show that this restriction continues to be as hard as the original problem, i.e.#P-hard. On the other hand, if G_M is restricted to be bimodal planar, or simultaneously planar and bipartite, then we show that computing $\mathsf{Perm}(M)$ is GapL-hard. This is the best lower bound possible, since in the next section we also show that in these cases we can also evaluate the permanent in GapL.

The results of this section can be summarized as follows:

Theorem 1. *The following problems are hard for* GapL *via* $\leq_m^{\log}$ *reductions.*

1. *DAG-WT-$s \rightsquigarrow t$-PATHS for planar graphs (total weight of all $s \rightsquigarrow t$ paths in a weighted directed acyclic graph DAG).*
2. Det*(M) for planar G_M (total signed weight of cycle covers in planar graph).*
3. Perm*(M) for planar bipartite G_M (total weight of cycle covers in G_M).*
4. Perm*(M) for planar bimodal G_M (total weight of cycle covers in G_M).*
5. Perm*(M) for planar bipartite H_M (total weight of perfect matchings in H_M).*

Further, computing Perm*(M) for planar G_M (total weight of cycle covers in planar graph) is hard for* #P.

We now sketch the proofs for each of these claims.

GapL $\leq_m^{\log}$ **Planar-DAG-WT-$s \rightsquigarrow t$-PATHS:** We start with the canonical GapL-complete problem Directed Path Difference (see for instance [26, 22]). The input is a directed graph G with special vertices s, t_+ and t_-, and the desired output $\#(G, s, t_+, t_-)$ is the difference in the number of $s \rightsquigarrow t_+$ paths and the number of $s \rightsquigarrow t_-$ paths. Without loss of generality, we can assume that (1) G is acyclic and layered (vertices appear in layers and all edges go from a layer to the next layer), (2) s is on the first layer and t_+ and t_- on the last layer, and all $s \rightsquigarrow t_+$ or $s \rightsquigarrow t_-$ paths are of even length, (3) all edges having weight 1, and (4) the number of vertices is odd.

We create a new vertex t and add edge $\langle t_+, t\rangle$ with weight 1, and edge $\langle t_-, t\rangle$ with weight -1, to get G_1. All $s \rightsquigarrow t$ paths are of odd length. The hard function is the total weight of all $s \rightsquigarrow t$ paths in G_1.

Now we planarize G_1 as follows: We draw G_1 in the plane, with edge crossings (as described in Proposition 1). We replace each crossing by the gadget shown alongside to get a planar graph G_2. Observe that for any vertices a, b in G_1, the weight of each $a \rightsquigarrow b$ path as well as the parity of the length of the path is preserved in G_2. Since G (and G_1) was bipartite, so is G_2. (Here bipartiteness is in the undirected sense: there are no undirected odd cycles.) Also, the embedding of G_2 we have is *upward planar*; it is planar and all edges are monotonic w.r.t. the x-coordinate [4]. In particular, this implies that the embedding of G_2 is bimodal. Without loss of generality, assume that G_2 has an odd number N of vertices.

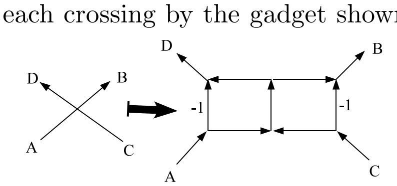

Fig. 1: Planarizing Gadget 1

We want to map paths in G_2 to (signed) cycle covers in a related graph. Toda [26] achieves this by subdividing every edge, adding self-loops everywhere except at s and then adding edge $\langle t, s\rangle$. We adapt this proof in two different ways.

[4] Using the techniques of Section 5, we can even ensure that G_2 is a *layered grid graph*.

$\mathsf{GapL} \leq_m^{\log}$ PLANAR 0-1 DET: The method of [26] does not eliminate negative weights. To handle this, we selectively subdivide only those edges with weight 1. Edges with weight -1 are not subdivided, but their weight is changed to 1. We can then show that this graph, say G_3, has the desired properties.

$\mathsf{GapL} \leq_m^{\log}$ {-1,0,1} BIPARTITE PLANAR BIMODAL DET/PERM: The above method loses bipartiteness not just because it adds self-loops, but also because of asymmetric subdivisions for weight 1 or -1. Instead, we can construct Split(G_2) and add to it edges $\langle v_{out}, v_{in} \rangle$ for each $v \notin \{s, t\}$, and the edge $\langle t, s \rangle$; all these edges have weight 1. Call this graph G_4; we can now show that it has the desired properties.

If A_3, A_4 are the adjacency matrices of G_3, G_4 respectively, then

$$\mathsf{Det}(A_4) = \mathsf{Perm}(A_4) = \mathsf{Det}(A_3) = \#(G_2, s, t) = \#(G_1, s, t) = \#(G, s, t_+, t_-)$$

GapL- $\leq_m^{\log}$ BIPARTITE PLANAR PERFECT MATCHINGS: Let G_5 be the undirected graph underlying Split(G_2); then G_5 is planar bipartite, and $s \rightsquigarrow t$ paths in G_2 are in 1-1 correspondence with perfect matchings in G_5 of the same weight. Thus the sum of the weights of the perfect matchings in G_5 is precisely $\#(G_2, s, t)$. (See [11, 18] for details.)

PERM $\leq_m^{\log}$ PLANAR PERM: We now show that computing $\mathsf{Perm}(M)$, when G_M is planar, is as hard as computing arbitrary permanents (i.e. $\#\mathsf{P}$-hard). Recall that $\mathsf{Perm}(M)$ computes the total weight of all cycle covers in G_M. Let N be the $n \times n$ matrix whose permanent we wish to compute. Consider the matrix $A = \begin{pmatrix} 0_n & N \\ I_n & 0_n \end{pmatrix}$ where I_n and 0_n denote the identity and the all-zeros matrices of size n. Clearly $\mathsf{Perm}(A) = \mathsf{Perm}(N)$. Consider a drawing of the directed bipartite graph G_A as described in Proposition 1.

As was done for the determinant, we replace each crossing with a planarity gadget so as to preserve the total weights of cycle covers. The planarity gadget used is shown alongside. Cycle covers using exactly one of the two edges AB or CD will now use the corresponding length 3 path $AXYB$ or $CYXD$. Cycle covers using neither of these edges will now use the 2-cycle XY. Cycle covers using both edges are essentially spliced; locally, we use instead the paths AXD and BYC.

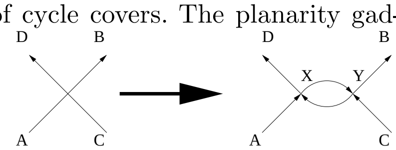

Fig. 2: Planarizing Gadget 2

Applying this planarity gadget to all crossings, we obtain a planar graph G_6 with adjacency matrix M. Since $\mathsf{Perm}(M) = \mathsf{Perm}(A) = \mathsf{Perm}(N)$, we have established the hardness of planar permanent.

Note that this planarity gadget preserves neither bipartiteness nor bimodality. This is not surprising, given the results of the next section.

4 Easy versions of Planar Permanent restrictions

We now show that certain planar restrictions of the permanent are significantly easier than #P, in fact, they are computable in GapL. We establish the following.

Theorem 2. *The following functions are computable in* GapL.

1. Perm(M) *for planar bipartite* G_M *(total weight of cycle covers in* G_M*).*
2. Perm(M) *for planar bimodal* G_M *(total weight of cycle covers in* G_M*).*
3. *Even-Odd Crossings Difference: The difference between the total weight of cycle covers with even number of crossings and the total weight of cycle covers with odd number of crossings, in a given plane drawing of a graph* G.

The proof of the first two results exploits the fact that finding the total weight of perfect matchings in planar graphs can be computed in GapL ([30, 21]).

Let $G_M = (V, E)$ be the given bipartite (directed) graph, with bipartition $X \dot{\cup} Y$. Let E_1 be those edges of E directed from X to Y, and E_2 be the remaining edges, and let $G_i = (V, E_i)$ for $i = 1, 2$ be planar bipartite undirected graphs. Then, with an appropriate renumbering of vertices (that can be computed in logspace since bipartite-testing is in L as a consequence of [24]), we have $M = \begin{pmatrix} 0_n & A_1 \\ A_2 & 0_n \end{pmatrix}$ where $H_{A_1} = G_1$ and $H_{A_2} = G_2$. (If G_M were undirected, we would have $A_1 = A_2^T$.) Clearly, $\mathsf{Perm}(M) = \mathsf{Perm}(A_1) \times \mathsf{Perm}(A_2)$. But $\mathsf{Perm}(A_i)$ equals the total weight of perfect matchings in the planar graph G_i, this can be computed in GapL.

If G_M is planar bimodal, then $\mathrm{Split}(G_M)$ is planar bipartite bimodal, and the cycle-covers in the two graphs are in bijection. So $\mathrm{Perm}(M)$ is the total weight of cycle covers in $\mathrm{Split}(G_M)$; we have just shown that this is in GapL.

The third result is really an exploration into how far planarizing gadgets can be pushed. If we can replace the crossings in a graph drawing by a gadget which preserves the weighted sum of cycle covers and also preserves bipartiteness or bimodality, then arbitrary pemanents would be expressible as planar bipartite permanents, implying the unlikely collapse of #P to GapL. This suggests that such gadgets are unlikely to exist.

However, we do have a bipartiteness-preserving gadget which reduces the Even-Odd Crossings Difference problem to cycle covers in planar graphs: Given a specific drawing of the graph, count the difference between the number of cycle covers with even number of crossings and the number of cycle covers with odd number of crossings. The gadget shown alongside will do the job. Now, if we start with a bipartite graph, then the resulting graph will be bipartite planar. So, for bipartite graphs, Even-Odd Crossings Difference can be computed in GapL.

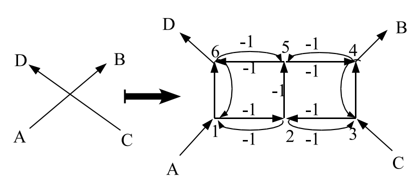

Fig. 3: Planarizing Gadget 3

5 Hardness of $\oplus$LGGR for $\oplus$L

Although the permanent is #P-hard, the permanent mod 2 equals the determinant mod 2 and is thus complete for $\oplus$L. A canonical $\oplus$L-complete problem is $\oplus$PATH-DAG: counting the number of $s \rightsquigarrow t$ paths, mod 2, in a directed acyclic graph (DAG). We show that this remains $\oplus$L hard (under $\leq_m^{\log}$-reductions) even if the DAG is planar, further, even if it is a layered grid graph. $\oplus$LGGR, referred to below, is layered grid graph reachability (LGGR) mod 2, that is, the problem of counting the number of $s \rightsquigarrow t$ paths mod 2 in a layered grid graph.

Theorem 3. $\oplus\mathsf{L} \leq_m^{log} \oplus\mathsf{LGGR}$

The result is significant because for the decision version (reachability in a DAG), the general case is NL-complete while its restriction to planar graphs is known to be in UL $\cap$ co-UL [9]. (Planar Directed Reachability PDR is known to be L-hard, and equivalent to reachability in grid graphs GGR [4], but its exact complexity is still unknown. Reachability in layered grid graphs LGGR is not even known to be L-hard. The complexity of various versions of grid graph reachability is investigated in [2].)

The following chain of reductions establishes the result.

$\oplus$PATH-DAG $\leq_m^{\log}$ $\oplus$PATH-PLANAR-DAG $\leq_m^{\log}$ $\oplus$PATH-x-MON-GG $\leq_m^{\log}$ $\oplus$PATH-LGG = $\oplus$LGGR

The first reduction considers a layered DAG (without loss of generality), draws it according to Proposition 1, and then uses the planarizing gadget of Figure 1, except that all edges have weight 1. This preserves the parity of the number of paths. From here, going to $\oplus$LGGR is achieved by using the grid-graph-embedding technique of [8, 10].

6 (Unique) Perfect Matchings in Planar Bipartite Graphs

We now investigate the complexity of checking existence and uniqueness of a perfect matching in a bipartite graph, B-PM and B-UPM, respectively when restricted to planar instances. Both B-PM and B-UPM are known to be NL-hard ([11, 18]), but B-UPM is believed to be easier since unlike B-PM, it is known to be in NC (in both $\mathsf{C_{=}L}$ and $\mathsf{NL}^{\oplus\mathsf{L}}$ [18]). We provide two further pieces of evidence that B-UPM may be easier by considering the planar restrictions of these problems, Pl-B-PM and Pl-B-UPM. Firstly, we show that while both are L-hard, Pl-B-PM is hard for Planar Directed Reachability PDR, whereas Pl-B-UPM is hard only for co-Layered Grid Graph Reachability co-LGGR. (It is known that PDR is equivalent to co-PDR and to its restriction Grid Graph Reachability GGR, [4]). The hardness of Pl-B-PM for PDR can be viewed as a planarization of the result "Reachability reduces to B-PM". We do not know how to planarize the result "co-Reachability reduces to bipartite-UPM" from [18]. Secondly, we obtain an upper bound of $\oplus$L for Pl-B-UPM. This can be viewed as a planarization of the result "B-UPM is in Reach$^{\oplus\mathsf{L}}$" from [18]: our algorithm is a $\mathsf{GGR}^{\oplus\mathsf{L}}$ algorithm, and since Section 5 shows that $\oplus$LGGR is hard for $\oplus$L, it is in fact in $\mathsf{GGR}^{\oplus\mathsf{LGGR}}$.

We note, however, that the complexity of LGGR (and co-LGGR) is an interesting question in its own right. It is not known whether it is in L, or L- hard, or reducible to its complement co-LGGR. However, its best known upper bound is the same as that for PDR, namely UL ∩ co-UL.

Also, analogous to the equivalence of PDR and GGR, we show that Pl-B-PM and Pl-B-UPM are equivalent to searching for or testing uniqueness of perfect matchings in grid graphs GGPM and GGUPM respectively.

We also consider the related problem of testing uniqueness of a minimum-weight perfect matching. In a bipartite graph with unary weights, this is known to be hard for NL and in $\mathsf{L}^{\mathsf{C_=L}}$ and $\mathsf{NL}^{\oplus\mathsf{L}}$ [18]. No better upper bound is known for the planar restriction, though the lower bound is also not known to hold. We show that GGR reduces to this planar restriction.

The results in this section can be summarized as follows. (See Figure 4. The pairs of dotted and dashed arrows show the planarizing results.)

Theorem 4. *1.* $(\mathsf{L} \cup co\text{-}\mathsf{LGGR}) \leq_{proj}$ Pl-B-UPM $\equiv_{proj}$ GGUPM $\in \oplus\mathsf{L}$

2. $(\mathsf{L} \cup \mathsf{GGR}) \leq_{proj}$ Pl-B-PM $\equiv_{proj}$ GGPM

3. Testing uniqueness of a min-weight perfect matching in a planar bipartite graph with unary weights is hard for GGR.

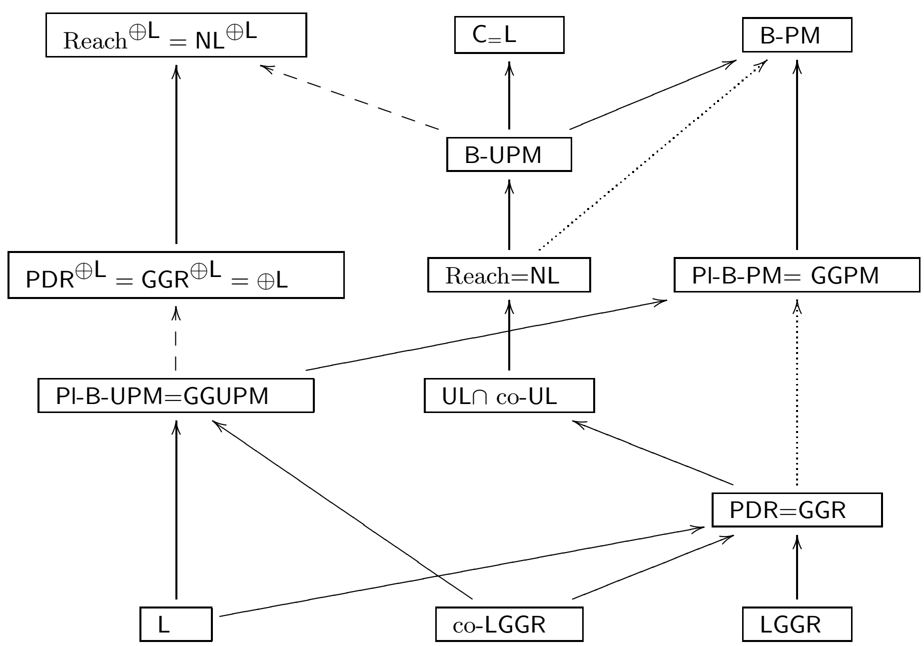

Fig. 4. Pl-B-UPM and Pl-B-PM and their relationships with other classes

L$\leq_{proj}$ Pl-B-UPM; L$\leq_{proj}$ Pl-B-PM: We start with the logspace-complete problem of determining whether there is an $s \rightsquigarrow t$ path in a directed forest G [12]. Given an instance (G, s, t), first construct its split graph G'. Then define H_1 to be the undirected version of G' and H_2 to be $H_1 \cup \{(s, t)\}$. Since G was a forest, H_1 and H_2 are clearly planar bipartite. Also their construction involves simple projections; it is FO-uniform.

Now, as in [11, 18], for every $s \rightsquigarrow t$ path in G, the alternate edges of the corresponding path in H, along with edges of the form (v_{in}, v_{out}) for vertices v not on the path, form a perfect matching in H_1 and H_2. H_1 has no other matching, H_2 has one more which is the added (s, t) edge along with all the edges of the form (v_{in}, v_{out}). Thus $H_1 \in$ Pl-B-PM if and only if $H_2 \in$ Pl-B-UPM if and only if (G, s, t) is not in Forest-Reachability.

co-LGGR$\leq_{proj}$ Pl-B-UPM; GGR$\leq_{proj}$ Pl-B-PM: This follows from carefully analysing the requirements in the above reduction, and some pre-processing.

Unique minimum weight Pl-B-UPM is hard for GGR: For the purpose of this section alone, the weight of a matching is the *sum* of its constituent edges.

Let (G, s, t) be the GGR instance; as discussed above, we can assume that G is bimodal and has s and t on the external face. We now assign weights to the edges of G according to the weighting scheme of [9] to get graph G'; this weighting scheme has the property that $s \rightsquigarrow_G t \iff s \rightsquigarrow_{G'} t \iff$ the minimum weight $s \rightsquigarrow_{G'} t$ path is unique. Now construct $H = \mathrm{Split}(G')$, copying the weight of an edge (u, v) in G' to the edge (u_{out}, v_{in}) of H and assigning weight zero toall the edges of the form (v_{in}, v_{out}). H is a planar bipartite graph and can be obtained via simple projections.

If $(G, s, t) \notin$ GGR, then it is easy to see that H has *no* perfect matching.

If $(G, s, t) \in$ GGR, then the unique minimum-weight path $\rho : s \rightsquigarrow_{G'}$ can be extended to a perfect matching in H $M_\rho = \{(u_{out}, v_{in}) \mid \langle u, v\rangle \in \rho_{G'}\} \cup \{(v_{in}, v_{out}) | v \in G' \text{and } v \notin \rho\}$ of the same weight. Since all (v_{in}, v_{out}) edges in H have weight 0, it is easy to see that this matching is the unique minimum-weight matching in H.

Pl-B-UPM$\leq_m^{\log}$ GGUPM; Pl-B-PM$\leq_m^{\log}$ GGPM: Both these results hold because there is a parsimonious (in the number of perfect matchings) reduction from planar bipartite graphs to grid graphs. This reduction is obtained by a slight modification of the grid graph embedding technique of [4], applied on an equivalent graph with maximum degree 3; the equivalent graph can also be obtained in logspace ([20]).

Pl-B-UPM is in $\oplus$L: In [18], an $\mathsf{NL}^{\oplus \mathsf{L}}$ algorithm for B-UPM is described. Given a bipartite graph G, it proceeds in two stages. In the first stage, an $\mathsf{L}^{\oplus \mathsf{L}}$ procedure either constructs some perfect matching M, or detects that G is not in B-UPM.

In the second stage, an NL procedure, with oracle access to M, verifies that M is indeed unique.

We show that for planar bipartite G, the second stage can be performed in $\mathsf{L}^{\mathsf{PDR}}$. Since PDR is known to be in $\mathsf{UL}\cap$co-UL[9] which is contained in $\oplus\mathsf{L}$, and since $\oplus\mathsf{L}^{\oplus\mathsf{L}} = \mathsf{L}^{\oplus\mathsf{L}} = \oplus\mathsf{L}$ ([17]), it then follows that $\mathsf{Pl\text{-}B\text{-}UPM}$ is in $\oplus\mathsf{L}$.

The key idea in obtaining the $\mathsf{L}^{\mathsf{PDR}}$ bound is the following: As described in [18], a given perfect matching M is unique in a bipartite graph G if and only if G has no alternating (with respect to M) cycles. We can consider an auxiliary graph H where an alternating path of length 2 in G, beginning with an M-edge, becomes a directed edge in H; then M is unique in G if H has no cycles. We show that H is planar. This implies that detecting cycles in H is in $\mathsf{L}^{\mathsf{PDR}}$.

References

1. E. Allender. Arithmetic circuits and counting complexity classes. In J. Krajicek, editor, *Complexity of Computations and Proofs*, Quaderni di Matematica Vol. 13, pages 33–72. Seconda Universita di Napoli, 2004. An earlier version appeared in the Complexity Theory Column, SIGACT News 28, 4 (Dec. 1997) pp. 2-15.
2. E. Allender, D. A. M. Barrington, T. Chakraborty, S. Datta, and S. Roy. Grid graph reachability problems. In *Proceedings of 21st IEEE Conference on Computational Complexity*, pages 299–313, 2006.
3. E. Allender, R. Beals, and M. Ogihara. The complexity of matrix rank and feasible systems of linear equations. *Computational Complexity*, 8(2):99–126, 1999.
4. E. Allender, S. Datta, and S. Roy. The directed planar reachability problem. In *Proc. 25th FSTTCS, LNCS vol. 3821*, pages 238–249, 2005.
5. E. Allender and M. Mahajan. The complexity of planarity testing. *Information and Computation*, 189(1):117–134, 2004.
6. E. Allender, K. Rheinhardt, and S. Zhou. Isolation, matching and counting: uniform and nonuniform upper bounds. *Journal of Computer and System Sciences*, 59:164–181, 1999.
7. D. Barrington. Bounded-width polynomial size branching programs recognize exactly those languages in NC^1. *Journal of Computer and System Sciences*, 38:150–164, 1989.
8. D. A. M. Barrington. Grid graph reachability problems. Talk presented at Dogstuhl Seminar on Complexity of Boolean funcions, Seminar Number 02121, 2002.
9. C. Bourke, R. Tewari, and N. V. Vinodchandran. Directed planar reachability is in unambiguous logspace. In *Proceedings of IEEE Conference on Computational Complexity CCC, to appear*, pages –, 2007.
10. T. Chakraborty and S. Datta. One-input-face MPCVP is hard for L, but in LogDCFL. In *Proc. of 26th FST TCS Conference, LNCS vol.*, pages –, 2006.
11. A. Chandra, L. Stockmeyer, and U. Vishkin. Constant depth reducibility. *SIAM Journal on Computing*, 13(2):423–439, 1984.
12. S. A. Cook and P. McKenzie. Problems complete for L. *Journal of Algorithms*, 8:385–394, 1987.
13. C. Damm. DET=$\mathrm{L}^{(\#\mathrm{L})}$. Technical Report Informatik–Preprint 8, Fachbereich Informatik der Humboldt–Universität zu Berlin, 1991.
14. A. L. Delcher and S. R. Kosaraju. An NC algorithm for evaluating monotone planar circuits. *SIAM Journal of Computing*, 24(2):369–375, 1995.

15. M. E. Dyer and A. M. Frieze. Planar 3DM is NP-complete. *J. Algorithms*, 7(2):174–184, 1986.
16. K. Hansen. Constant width planar computation characterizes ACC^0. In *Proc. 21st STACS, LNCS vol. 2996*, pages 44–55, 2004.
17. U. Hertrampf, S. Reith, and H. Vollmer. A note on closure properties of logspace MOD classes. *Information Processing Letters*, 75(3):91–93, 2000.
18. T. M. Hoang, T. Thierauf, and M. Mahajan. On the bipartite unique perfect matching problem. In *Proc. 33rd International Colloquium on Automata, Languages and Programming (ICALP)*, pages 453–464. Springer, 2006. LNCS 4051.
19. H. B. Hunt, III, M. V. Marathe, V. Radhakrishnan, and R. E. Stearns. The complexity of planar counting problems. *SIAM Journal on Computing*, 27(4):1142–1167, 1998.
20. R. Kulkarni and M. Mahajan. Seeking a vertex of the planar matching polytope in nc. In *Proceedings of the 12th European Symposium on Algorithms ESA, LNCS vol. 3221*, pages 472–483. Springer, 2004.
21. M. Mahajan, P. R. Subramanya, and V. Vinay. The combinatorial approach yields an NC algorithm for computing Pfaffians. *Discrete Applied Mathematics*, 143(1-3):1–16, September 2004.
22. M. Mahajan and V. Vinay. Determinant: combinatorics, algorithms, complexity. *Chicago Journal of Theoretical Computer Science* `http://www.cs.uchicago.edu/publications/cjtcs`, 1997:5, Dec 1997.
23. B. Mohar and C. Thomassen. *Graphs on Surfaces*. John Hopkins University Press, Maryland, 2001.
24. O. Reingold. Undirected *st*-connectivity in logspace. In *Proc. 37th STOC*, pages 376–385, 2005.
25. K. Reinhardt and E. Allender. Making nondeterminism unambiguous. *SIAM J. Comp.*, 29:1118–1131, 2000.
26. S. Toda. Counting problems computationally equivalent to the determinant. Technical Report CSIM 91-07, Dept of Comp Sc & Information Mathematics, Univ of Electro-Communications, Chofu-shi, Tokyo, 1991.
27. S. Vadhan. The complexity of counting in sparse, regular, and planar graphs. *SIAM Journal on Computing*, 31(2):398–427, 2001.
28. L. G. Valiant. The complexity of computing the permanent. *Theoretical Computer Science*, 8:189–201, 1979.
29. L. G. Valiant. Why is boolean complexity theory difficult? In M. S. Paterson, editor, *Boolean Function Complexity*. Cambridge University Press, 1992. London Mathematical Society Lecture Notes Series 169.
30. V. Vazirani. NC algorithms for computing the number of perfect matchings in $K_{3,3}$-free graphs and related problems. *Information and Computation*, 80(2):152–164, 1989.
31. V. Vinay. *Semi-unboundedness and complexity classes*. PhD thesis, Indian Institute of Science, Bangalore, July 1991.
32. M. Xia and W. Zhao. #3-regular bipartite planar vertex cover is #P-complete. In *TAMC*, pages 356–364, 2006.
33. H. Yang. An NC algorithm for the general planar monotone circuit value problem. In *Proceedings of 3rd IEEE Symposium on Parallel and Distributed Processing*, pages 196–203, 1991.