

Parametric Shortest Paths in Planar Graphs

Kshitij Gajjar and Jaikumar Radhakrishnan

Tata Institute of Fundamental Research, Mumbai, India

Email: {kshitij.gajjar, jaikumar}@tifr.res.in

Abstract—We construct a family of planar graphs $\{G_n\}_{n \geq 4}$, where G_n has n vertices including a source vertex s , a sink vertex t , and edge weights that change linearly with a parameter λ such that, as λ varies in $(-\infty, +\infty)$, the piece-wise linear cost of the shortest path from s to t has $n^{\Omega(\log n)}$ pieces. This shows that lower bounds obtained by Carstensen (1983) and Mulmuley & Shah (2001) for general graphs also hold for planar graphs, refuting a conjecture of Nikolova (2009).

Gusfield (1980) and Dean (2009) showed that the number of pieces for every n -vertex graph with linear edge weights is $n^{\log n + O(1)}$. We generalize this result in two ways. (i) If the edge weights vary as a polynomial of degree at most d , then the number of pieces is $n^{\log n + (\alpha(n) + O(1))^d}$, where $\alpha(n)$ is the inverse Ackermann function. (ii) If the edge weights are linear forms of three parameters, then the number of pieces, appropriately defined for \mathbb{R}^3 , is $n^{(\log n)^2 + O(\log n)}$.

Index Terms—parametric complexity; shortest paths; planar graphs; alternation-free sequences

I. INTRODUCTION

We consider the following *parametric shortest path problem* on graphs. The input is a directed acyclic graph with two special vertices s and t . The edges have weights that vary linearly with a real-valued parameter λ , that is, the weight of each edge e is a function of the form $w_e(\lambda) = a_e\lambda + b_e$, for some real numbers a_e and b_e (Figure 1). The cost of an s - t path P is the sum of the weights of the edges on it; therefore this cost is also a linear function of λ of the form $\mathcal{C}(P)(\lambda) = \sum_{e \in P} a_e\lambda + \sum_{e \in P} b_e$. The cost of the shortest s - t path is then given by

$$\mathcal{C}(\lambda) = \min_P \mathcal{C}(P)(\lambda),$$

where P ranges over all s - t paths; this function is the piece-wise linear lower envelope (Figure 2) of the linear costs provided by the s - t paths. The main object of our investigation is the number of pieces in this envelope. This quantity is of interest in several applications; in particular, determining this quantity for *planar graphs* has been a subject of several studies.

Let the *parametric shortest path complexity*, denoted by $\varphi(n, \beta(n))$, be the maximum possible number of pieces in $\mathcal{C}(\lambda)$ for graphs on n vertices, where the bit lengths of the coefficients in the weights of the edges are bounded by $\beta(n)$. Let $\varphi^{\text{pl}}(n, \beta(n))$ be the parametric complexity when the graphs are restricted to be planar. We show the following.

Theorem 1 (Main result, lower bound on the parametric complexity of planar graphs).

$$\varphi^{\text{pl}}(n, O((\log n)^3)) = n^{\Omega(\log n)}.$$




Figure 1: A planar directed acyclic graph G (horizontal edges go rightwards and vertical edges go upwards). The weights of the edges of G are linear functions of λ .




Figure 2: A plot of λ on the X-axis versus the (linear) costs of the 6 different s - t paths in the graph G on the Y-axis. The piece-wise linear lower envelope of this plot has 4 pieces.

Similar results were known for general graphs. Carstensen ([1], [2]) showed that $\varphi(n, \infty) = n^{\Omega(\log n)}$; her result was simplified and extended by Mulmuley & Shah [3], who showed that $\varphi(n, O((\log n)^3)) = n^{\Omega(\log n)}$. For *planar graphs*, Nikolova [4, Conjecture 6.1.6] conjectured that the complexity is bounded by a polynomial in n , that is, $\varphi^{\text{pl}}(n, \infty) = n^{O(1)}$. Our main result provides a strong (with bit length $O((\log n)^3)$) refutation of this conjecture.

The above lower bounds are tight. Carstensen [2, Page 100] presented a matching upper bound, $\varphi(n, \infty) = n^{\log n + O(1)}$, which she attributed to Daniel Gusfield ([5], [6]). A sim-

ilar argument, attributed to Brian Dean, was presented by Nikolova [4, Page 86]. We generalize these upper bounds in two ways based on how the edge weights w_e vary: (i) $w_e(\lambda)$ is a polynomial of degree at most d in $\lambda \in \mathbb{R}$, and (ii) $w_e(\lambda_1, \lambda_2, \lambda_3) = a_e \lambda_1 + b_e \lambda_2 + c_e \lambda_3$, where $(\lambda_1, \lambda_2, \lambda_3) \in \mathbb{R}^3$. A slightly different definition of φ is required for these generalizations.

Consider a graph G whose edge weights depend on a parameter λ taking values in a set Λ . As λ varies, the minimum cost s - t path might vary. We say that a set \mathcal{P} of s - t paths covers G if for every value $\lambda \in \Lambda$, the set \mathcal{P} contains a minimum cost s - t path of G . We define the parametric shortest path complexity of G as

$$\varphi(G) = \min_{\mathcal{P}: \mathcal{P} \text{ covers } G} |\mathcal{P}|.$$

For the generalization (i), let $\varphi^{\deg(d)}(n)$ be the maximum value of $\varphi(G)$ as G varies over n -vertex graphs whose edge weights are polynomials of degree at most d in a parameter $\lambda \in \mathbb{R}$. We use the inverse Ackermann function (Definition 30) to upper bound $\varphi^{\deg(d)}(n)$.

Theorem 2 (Upper bound with univariate polynomial edge weights).

$$\varphi^{\deg(d)}(n) = n^{\log n + (\alpha(n) + O(1))^d},$$

where $\alpha(n)$ is the extremely slow growing inverse Ackermann function.

For the generalization (ii), let $\varphi^{\text{lin}(k)}(n)$ be the maximum value of $\varphi(G)$ as G varies over n -vertex graphs whose edge weights have the form $w_e(\vec{\lambda}) = \vec{a}_e \cdot \vec{\lambda}$, where $\vec{a}_e \in \mathbb{R}^k$ and $\vec{\lambda} \in \mathbb{R}^k$.

Theorem 3 (Upper bound with three-parameter linear edge weights).

$$\varphi^{\text{lin}(3)}(n) = n^{(\log n)^2 + O(\log n)}.$$

Remarks: (i) Upper bounds in Theorem 2 and Theorem 3 grow only moderately (for small d). (ii) Theorem 3 leads to the natural question whether similar bounds can be shown for $\varphi^{\text{lin}(k)}(n)$ in general; unfortunately, our proof method fails when $k > 3$. (iii) A bound of the form $\varphi^{\text{lin}(2)}(n) = n^{\log n + O(1)}$ can be derived using our method; this implies Gusfield's bound cited above.

A. Significance of the Main Result

In this section, we present some consequences of Theorem 1.

a) PRAM lower bounds: From their result (that is, $\varphi(n, O((\log n)^3)) = n^{\Omega(\log n)}$), Mulmuley & Shah [3] derived a lower bound on the running time of unbounded fan-in PRAMs with bit operations with a small number of processors solving the shortest path problem. Theorem 1 allows us to make a similar claim for planar graphs (see Section VIII for a discussion on this).

Claim 4. *There exist constants $\alpha > 0, \epsilon > 0$, and an explicitly described family of weighted planar graphs $\{G_n\}$ (G_n has n*

vertices, and the edge weights of G_n are $O((\log n)^3)$ bits long), such that for infinitely many n , every unbounded fan-in PRAM algorithm (without bit operations) with at most n^α processors requires at least $\epsilon \log n$ steps to compute the shortest s - t path in G_n .

b) Weighted graph matching: Mulmuley & Shah observed that their result for the shortest path problem yields the same lower bound for the WEIGHTED GRAPH MATCHING problem [3, Corollary 1.1]. Our result extends this observation to planar graphs (see Section IX for a proof of this reduction). Many graph problems are easier to solve for planar graphs than for general graphs; in particular, we note the NC algorithm for counting perfect matchings based on the work of Kasteleyn [7] and Csanky [8], and its remarkable recent application by Anari & Vazirani [9] (see also [10]) to find perfect matchings in planar graphs. It is interesting that the lower bound for the WEIGHTED GRAPH MATCHING problem derived by Mulmuley & Shah continues to hold even when the input is restricted to be planar.

c) Treewidth: Planar graphs have high (superpolynomial) parametric complexity (Theorem 1). It is thus natural to ask what graph classes might have small parametric complexity. Since every planar graph can be embedded in a grid graph with a marginal increase in size, our lower bound holds for $n \times n$ grid graphs as well. We also explore the parametric complexity of $k \times n$ grid graphs when $k \ll n$ (see Section IV). Due to a result of Chekuri & Chuzhoy [11], every graph of treewidth k has an $\Omega(k^{1/99}) \times \Omega(k^{1/99})$ grid minor. Thus, for large enough k , every graph of treewidth k has parametric complexity $k^{\Omega(\log k)}$. In particular, if the graph class has n -vertex graphs whose treewidth grows as $k(n) = \exp(\omega(\sqrt{\log n}))$, then its parametric complexity grows superpolynomially. On the other hand, our construction shows that for every $k(n) = \omega(\log n)$, there are planar graphs with treewidth $k(n)$ and parametric complexity $n^{\omega(1)}$; in the reverse direction, it can be shown that n -vertex graphs of treewidth k have parametric complexity $n^{O(k)}$ [12].

d) Minimum weight s - t cut: Our planar graphs are built such that s and t lie on the same face when the graph is drawn on a plane. By appealing to the planar dual of our graph, we conclude that the parametric complexity of the (s, t) -cut problem in planar graphs is also $n^{\Omega(\log n)}$.

e) Undirected graphs: Our construction yields a directed graph, but with a slight modification (by increasing all edge costs uniformly), we obtain an undirected graph with the same number of pieces. Thus our $n^{\Omega(\log n)}$ lower bound holds for undirected planar graphs as well.

f) Optimization problems: Parametric shortest paths have been studied extensively in the optimization literature because of their close connection with several other problems. We briefly mention four.

- Nikolova, Kelner, Brand & Mitzenmacher [13] consider a stochastic optimization problem on graphs whose edge weights represent random Gaussian variables and where one is required to determine the s - t path whose total

cost is most likely to be below a specified threshold (the deadline). They provide an $n^{O(\log n)}$ time algorithm for the problem for general graphs, and suggest that when restricted to planar graphs their algorithm might run in polynomial time because the number of extreme points of the *shadow dominant* (a notion closely related to parametric shortest path complexity) is likely to be polynomially (perhaps even linearly) bounded. Our result unfortunately belies this hope.

- Correa, Harks, Kreuzen & Matuschke [14] study the problem of *fare evasion in transit networks*, and consider strategies based on random checks for the service providers, and the response of the users to such strategies. For one of the problems, referred to as the non-adaptive followers' minimization problem, they devise an algorithm based on the parametric shortest path problem, and point out that their algorithm would run in polynomial time on planar graphs if Nikolova's conjecture were to hold.
- Chakraborty, Fischer, Lachish & Yuster [15] provide two-phase algorithms for the parametric shortest path problem, where the first stage does preprocessing after which an advice is stored in memory so that the algorithm can answer queries efficiently thereafter. A natural application for such an algorithm is traffic networks. Since traffic networks tend to be planar, a good upper bound on the parametric complexity of planar graphs would have allowed for substantial savings in space.
- We also mention work on a closely related problem. Erickson [16] reformulates an $O(n \log n)$ time algorithm of Borradaile & Klein [17] for max-flows in planar graphs by considering parametric shortest path trees (see Karp & Orlin [18]) in the dual graph. He shows that the tree can undergo only a limited number of changes. However, in Erickson's setting, the coefficient of λ in the edge weights is always -1 . He also points out that a similar approach for max-flows in graphs drawn on a torus fails to yield an efficient algorithm because the tree might undergo $\Omega(n^2)$ changes.

B. Overview of our Proof Techniques

We recall two earlier efforts aimed towards resolving Nikolova's conjecture. In her PhD thesis, Nikolova [4] considers planar embeddings of planar graphs, and shows that the edges can always be assigned weights in such a way that the number of pieces (in the piece-wise linear plot of the shortest paths) is at least the number of faces in the embedding. Note, however, that the number of pieces in the n -vertex planar graphs constructed using this approach is at most $2n$. We are aware of only one work that establishes a better upper bound for a family of planar graphs: Correa *et al.* [14] observe that for series parallel graphs, Nikolova's conjecture is true; the parametric complexity of series-parallel graphs is in fact $O(n)$.

Proof techniques for the main result: It is instructive¹ to

¹As perhaps many others did before us, we initially believed that Nikolova's conjecture was true and tried to prove it.

briefly review the upper bound arguments of Gusfield and Dean with the hope of tightening them in the setting of planar graphs. Let $G[n, m]$ denote a directed acyclic graph G with vertices s and t that has m layers of n vertices each in between s and t . Fix a numbering of the vertices $(1, 2, \dots, n)$ in each layer. These arguments are based on the following observations. Let us assume that the shortest s - t path is constructed in such a way that starting from s we always move to the neighbour with the shortest distance to t , choosing the neighbour having the smallest number when there is a tie. Let (p_1, p_2, \dots, p_T) be the sequence of shortest paths corresponding to the lower envelope, where each path p_i is constructed in this fashion. This sequence of paths has the following *alternation-free property* (called *expiration property* by Nikolova [4]). For a path p , and vertices u and v that appear on it in that order, let $p[u : v]$ be the subpath of p that connects u to v .

Proposition 5 (Alternation-free property, expiration property). *Suppose vertices u and v both appear on the three paths p_i , p_j and p_k in the sequence (p_1, p_2, \dots, p_T) , where $1 \leq i < j < k \leq T$. Furthermore, suppose $q = p_i[u : v] = p_k[u : v]$. Then, $p_j[u : v] = q$.*

This alternation-free property is important because the length of a longest alternation-free sequence of paths in n -vertex planar graphs is an upper bound on $\varphi^{\text{pl}}(n, \infty)$.

Theorem 6 (Alternation-free sequences of paths in layered graphs). *Let $G[n, m]$ be the layered graph with $m = 2^k$ layers, each layer consisting of n vertices (s is connected to all vertices in the first layer and t is connected to all vertices in the last layer). Let $f(n, m)$ be the length of a longest alternation-free sequence paths in $G[n, m]$; let $f^{\text{pl}}(n, m)$ be the length of a longest alternation-free sequence of paths in a planar subgraph of $G[n, m]$. Then, we have the following.*

$$(i) f(n, 2^k) \geq n^k, (ii) f^{\text{pl}}(n, (n-1)2^k) \geq n^k.$$

Using the alternation-free property, observe that $f(n, 1) = n$ and $f(n, 2^k - 1) \leq 2nf(n, 2^{k-1} - 1)$, which yields $f(n, 2^k - 1) \leq \frac{1}{2}(2n)^k$, implying that $\varphi(n, \infty) = n^{O(\log n)}$. Note that the *non-planar graphs* with high parametric shortest path complexity constructed by Carstensen [2] and Mulumuley & Shah [3] imply that $f(n, n) \geq n^{\delta \log n}$ (for some $0 < \delta < 1$). In Subsection VII-B, we present a construction which shows that $f(n, 2^k) \geq n^k$. Thus, we have $n^k \leq f(n, 2^k) \leq \frac{1}{2}(2n)^k$. More crucially, our construction can be adapted to *planar graphs*.

In Subsection VII-C, we present the construction for planar graphs in detail. This shows that just using the alternation-free property alone, we cannot hope to obtain significantly better upper bounds on $\varphi^{\text{pl}}(n, \infty)$. While this construction provides some evidence against Nikolova's conjecture, it does not immediately refute it. In fact, there exist examples of alternation-free sequences of paths in planar graphs that do not arise as parametric shortest paths. Kuchlbauer [19, Example 3.11] presents a planar graph that admits an *infeasible* alternation-

free sequence with 10 paths; that is, no assignment of linear functions to the edges can realize this sequence of 10 paths as shortest paths.

Our refutation of Nikolova’s conjecture is based on the Mulmuley-Shah construction [3]. Their construction uses an intricate inductive argument involving the composition of dense bipartite graphs. These bipartite graphs contain large complete bipartite graphs, and are therefore far from planar. We show that, nevertheless, these non-planar bipartite graphs can be simulated by a planar gadget, where each edge is replaced by a path consisting of up to n^2 edges and the original weight is carefully distributed among these edges. For this we introduce two ideas. First, staying with the original non-planar construction, we modify the edge weights so that they vary in a structured way. Second, we imagine that the original bipartite graph is drawn on a plane by connecting dots using straight lines, a new vertex arising whenever two straight lines intersect. This results in several new vertices, and spurious paths that do not correspond to any edge of the original bipartite graph. However, the costs of the new edges are so assigned that these spurious paths have much higher costs than the direct path corresponding to the edge in the original bipartite graph. We devote Section II to the construction of this gadget.

The main technique in our construction goes back to Carstensen’s work. Our planarization is straightforward in hindsight. The reasons this was not observed before are perhaps the following: (i) the earlier recursive constructions even for general graphs are complicated and not easy to take apart and examine closely (in particular, the Mulmuley-Shah paper is rather cryptic and has errors that throw the reader off); (ii) simple methods of constructing planar graphs with many pieces (in the piece-wise linear plot) tend to navigate around regions in the planar drawing one at a time, somehow (mis)leading one to believe that the limited number of planar regions ought to impose a polynomial upper bound on the number of pieces. See Section III for a detailed proof.

Proof techniques for the upper bounds: Recall the upper bound arguments for alternation-free paths leading to the recurrence stated after Proposition 5. When edge weights vary as degree d polynomials and not just linearly, paths are no longer strictly alternation-free; rather, two paths could alternate up to d times. The complication arising out of this is related to *Davenport-Schinzel sequences*, which have been studied extensively in the discrete geometry literature. The existing upper bounds on the length of Davenport-Schinzel sequences can be combined with the approach of Gusfield and Dean to yield Theorem 2. See Section V for a detailed proof.

However, when edge weights have the form $a\lambda_1 + b\lambda_2 + c\lambda_3$, our proof techniques depart substantially from the arguments used by Gusfield (which were adapted to univariate polynomials to obtain Theorem 2). Although Gusfield’s bound is stated for edge weights of the form $a\lambda + b$, essentially the same upper bound holds when the edge weights have the form $a\lambda_1 + b\lambda_2$. (This can be seen, for example, by dividing all costs uniformly by λ_2 .) In the three-parameter setting, it is not




Figure 3: $G_{m,n,B}$ is obtained by composing $G_{m-1,n,B}$, $G_{m-1,n,B}^{\text{rev}}$, the linking gadget, and $G_{m-1,n,B+n}$

clear how one can impose a meaningful linear order on the set \mathbb{R}^3 , and invoke combinatorial notions such as alternation-free sequences. Instead, we approach the problem geometrically.

Note that the cost of an s - t path P has the form $a_P\lambda_1 + b_P\lambda_2 + c_P\lambda_3$, and may be viewed as a vector (a_P, b_P, c_P) in \mathbb{R}^3 . Consider the convex hull of these path vectors. The crucial observation is that for each $\vec{\lambda} = (\lambda_1, \lambda_2, \lambda_3) \in \mathbb{R}^3$, at least one of the vertices (vertex here means 0-dimensional face) of the convex hull corresponds to a minimum cost path. Thus, we need an upper bound on the number of vertices of the convex hull. From here on, our argument is similar to Gusfield’s but is carried out in the language of polytopes. The key non-trivial step in the analysis of the number of vertices in the polytope arises when two graphs are placed in series. This is addressed by bounding the number of vertices of the *Minkowski sum* of the polytopes of the constituent graphs. A detailed proof of Theorem 3 is presented in Section VI.

We now briefly point out that there is an obstruction to a simple reduction to the case of two variables. Suppose we set λ_3 to 1. Then we are left with a situation where the cost of each s - t path is a plane in \mathbb{R}^3 . The shortest path function $C(G)(\lambda_1, \lambda_2)$ is then a dome-like structure, a polyhedron representing the lower envelope of the planes corresponding to the various s - t paths. We wish to bound the number of faces of this polyhedron. Note that every planar cross-section of this polyhedron is a piece-wise linear function in \mathbb{R}^2 , which by Gusfield’s bound has $n^{\log n + O(1)}$ pieces. A conjecture of Shephard [20] states that the number of faces in the polyhedron is bounded by a polynomial in the maximum number of pieces in such a cross-section. Unfortunately, this conjecture is false: there are polyhedrons with n faces for which the number of pieces in every planar cross-section is at most $O(\log n / \log \log n)$ (see [21], [22]).

II. THE PLANAR CONSTRUCTION

In this section, we construct a planar gadget which will be used to construct planar graphs with high shortest path complexity. Our construction closely follows the construction of Mulmuley & Shah [3], which in turn was based on the

construction of Carstensen [2]. These earlier constructions (and ours) proceed by induction, where we begin with a small base graph, and at each induction step, we increase the number of vertices by a constant factor and the number of pieces in the lower envelope by a factor n . After m steps of induction, we obtain a graph with $\text{poly}(n) \cdot \exp(O(m))$ vertices and n^m pieces. Figure 3 in [23] illustrates this assembly for $m = 3, n = 3$, following the template of Figure 3. We will explain this construction in detail later. The edge weights in the constituent graphs in Figure 3 are carefully chosen, but are not important to our top-level view. The only new component added in each level of induction is the part labelled LINK.

Our first observation is that the edge weights used by Mulmuley & Shah in LINK can be modified so that they have a regular form. Our second observation is that with the modified edge weights, LINK, which is a dense bipartite graph, can be simulated by a planar gadget.

In the following sections, we provide detailed justification for the two contributions outlined above. In Section III, we show that the new edge weights in LINK also result in a large number of pieces in the lower envelope. In Subsection II-A, we show that the non-planar graph G^{np1} can be simulated by a suitable weighted planar graph; this step, which is at the core of our contribution, has a simple implementation with an appealing proof of correctness.

A. The Linking Gadget

A linking gadget LINK, denoted by $\mathbf{L}(B, n)$, is a bipartite graph $G(U, V, E, (w_e : e \in E))$ with $U = \{0, 1, \dots, B-1\}$, $V = \{0, 1, \dots, B+n-1\}$, $E = \{(b, b+r) : b \in U, r = 0, 1, \dots, n\}$ (see Figure 4). In this graph the cost of the shortest path from vertex b to vertex j is precisely $w_{(b,j)}$ (we often write $w_{b,j}$ instead). We would like to obtain a directed planar simulation of this behaviour.

Let G^{pl} be the directed graph drawn on a planar strip in \mathbb{Q}^2 given by $[0, 1] \times [0, 2n-2]$; the vertices of G^{pl} include the sets of points $\{0\} \times U$ and $\{1\} \times V$; the rest of the graph is obtained as follows. We draw the line segments $\ell_{(b,j)}$ joining $(0, b)$ to $(1, j)$ whenever $(b, j) \in E(G)$, and include all intersection points of such segments in the vertex set of G^{pl} (see Figure 5). The edge (u, v) is in G^{pl} if v immediately follows u on some line segment ℓ_e . The edge weight w_e of the edge $e \in E(G)$ is distributed uniformly among the various edges of G^{pl} that arise out of e . Suppose the vertices $u = (u_x, u_y)$ and $v = (v_x, v_y)$ appear consecutively on ℓ_e (note $v_x > u_x, v_y \geq u_y$); then $w_{u,v} = w_e \cdot (v_x - u_x)$.

This completes the description of the weighted planarization G^{pl} (also denoted by $\mathbf{L}^{\text{pl}}(B, n)$) of G . Note that G^{pl} has $\text{poly}(B, n)$ vertices. The locations of the vertices in this special planar embedding of G^{pl} are not essential for our construction. However, one feature of this embedding is useful in our proof. A vertex is placed at a point of intersection of two lines of the form $Y = m_1X + c_1$ and $Y = m_2X + c_2$; so its x -coordinate, namely $(c_2 - c_1)/(m_1 - m_2)$ can be written as a fraction with denominator at most n , which leads to the following observation.




Figure 4: The non-planar LINK gadget $\mathbf{L}(B, n)$ for $B = 4, n = 3$




Figure 5: The planarized LINK gadget $\mathbf{L}^{\text{pl}}(B, n)$ for $B = 4, n = 3$

Fact 7. *The total horizontal distance traversed by an edge $((x_1, y_1), (x_2, y_2))$ of G^{pl} (which is $x_2 - x_1$) can be expressed as a non-zero fraction with denominator at most n^2 .*

We will use this observation later (Claim 10). Before that, we need to formalize what it means for G^{pl} to mimic G .

Definition 8. *We say that G^{pl} faithfully simulates G if for all $(b, j) \in U \times V$:*

- (i) *if $b \leq j \leq b + n$, then the edges arising from the line segment $(0, b)$ to $(1, j)$ form the unique shortest path from $(0, b)$ to $(1, j)$ in G^{pl} ,*
- (ii) *if $b \leq j \leq b + n$, then the cost of the shortest path from $(0, b)$ to $(1, j)$ in G^{pl} is precisely $w_{b,j}$, and the cost of every other path from $(0, b)$ to $(1, j)$ is at least $w_{b,j} + 1$,*

and

- (iii) if $j < b$ or $j > b + n$, then there is no path from $(0, b)$ to $(1, j)$ in G^{pl} .

In spirit, this definition says that shortest paths in G^{pl} should look like edges of G . \diamond

Lemma 9. Suppose $J : E(G) \rightarrow \mathbb{Z}$, and K and L are constants such that

$$K \geq n^2 \left(1 + 2 \max_{e \in E(G)} |J(e)| \right).$$

Consider a graph $G(U, V, E, w)$ of the form described above with edge weights

$$w_{b, b+r} = J(b, b+r) + K \left(\frac{r(r+1)}{2} \right) + Lr\lambda,$$

where $0 \leq b \leq B-1$ and $0 \leq r \leq n$. Then G^{pl} faithfully simulates G . Note that the rate of change of $w_{b, b+r}$ w.r.t. λ is Lr , where r is the slope of the line segment $\ell_{(b, b+r)}$ w.r.t. the X -axis.

Proof. Consider vertices $b \in U$ and $j \in V$ such that $0 \leq b \leq j \leq b+n$. Consider the path P in G^{pl} that takes edges along the line segment $\ell_{(b, j)}$. This path has cost $w_{b, j}$. We will show that all other paths from $(0, b)$ to $(1, j)$ have strictly greater cost. Let $r = j - b$ be the slope of the line segment $\ell_{(b, j)}$. Suppose Q is another path in G^{pl} from $(0, b)$ to $(1, j)$.

We make the following claim.

Claim 10. $\mathcal{C}(Q) - \mathcal{C}(P) \geq n^{-2}K - 2 \max_{e \in E(G)} |J(e)|$.

Proof of claim. Let Q consist of vertices $q_0 = (x_0, y_0), q_1 = (x_1, y_1), q_2 = (x_2, y_2), \dots, q_t = (x_t, y_t)$, where $(x_0, y_0) = (0, b)$ and $(x_t, y_t) = (1, j)$ for some b, j . For $i = 1, 2, \dots, t$, let $r_i = (y_i - y_{i-1}) / (x_i - x_{i-1})$ denote the slope of the edge (q_{i-1}, q_i) ; let $\rho_i = x_i - x_{i-1}$. Then for $i \in \{1, 2, \dots, t\}$, we have

$$\rho_i \geq n^{-2}; \quad (\text{using Fact 7})$$

$$r_i \in \{0, 1, \dots, n\};$$

$$r = \sum_{i=1}^t \rho_i r_i = j - b;$$

$$w_{q_{i-1}, q_i} = \left(J(y, y+r_i) + K \left(\frac{r_i(r_i+1)}{2} \right) + Lr_i\lambda \right) \rho_i,$$

where $y = y_i - r_i x_i$. Since $0 \leq \rho_i \leq 1$ and $\sum_{i=1}^t \rho_i = 1$, we may define a random variable \mathbf{i} , that takes the value $i \in \{0, 1, \dots, n\}$ with probability ρ_i . With this notation, we have $\mathbb{E}[r_{\mathbf{i}}] = r$.

Alternative view: This paragraph provides a physics perspective to the proof. The calculations can be carried out without reading this paragraph. We may view the path as a height versus time graph of a moving particle, with the X -axis representing time and the Y -axis representing height. Then, $\mathbb{E}[r_{\mathbf{i}}] = r = j - b$ corresponds to the fact that the particle underwent a displacement of r units in one unit of time. Note that this claim holds independent of the path taken

by the particle. As a result, for the purpose of comparing costs of paths, we can ignore the terms containing λ . Let us now proceed with the calculations.

$$\begin{aligned} \mathcal{C}(Q) - \mathcal{C}(P) &\geq -2 \max_e |J(e)| + K \left(\mathbb{E} \left[\frac{r_{\mathbf{i}}^2}{2} \right] - \frac{r^2}{2} \right) \\ &\quad + \left(\frac{K + 2L\lambda}{2} \right) (\mathbb{E}[r_{\mathbf{i}}] - r) \\ &\geq -2 \max_e |J(e)| + K \left(\mathbb{E} \left[\frac{r_{\mathbf{i}}^2}{2} \right] - \frac{r^2}{2} \right) \\ &\geq -2 \max_e |J(e)| + \frac{K}{2} \text{var}[r_{\mathbf{i}}]. \end{aligned}$$

In these calculations, we used the fact that $\mathbb{E}[r_{\mathbf{i}}] = r$. We now show a lower bound for $\text{var}[r_{\mathbf{i}}]$. Since Q deviates from P , it has at least two edges whose slopes, say r_{i_1} and r_{i_2} , differ from r (by at least 1). Then,

$$\text{var}[r_{\mathbf{i}}] \geq \rho_{i_1} (r_{i_1} - r)^2 + \rho_{i_2} (r_{i_2} - r)^2 \geq 2n^{-2}.$$

This establishes Claim 10. \square

The assumption on K then implies that P is the unique shortest path from $(0, b)$ to $(1, j)$, and the cost of every other path Q from $(0, b)$ to $(1, j)$ is at least $w_{b, j} + 1$. This proves (i) and (ii). Finally, (iii) holds because every edge in G^{pl} corresponds to a line segment with slope at least 0 and at most n . This completes the proof of Lemma 9. \square

III. PROOF OF THE MAIN RESULT

Theorem 1 (Main result, lower bound on the parametric complexity of planar graphs).

$$\varphi^{\text{pl}}(n, O((\log n)^3)) = n^{\Omega(\log n)}.$$

Our proof is constructive. Throughout this section, $n \in \mathbb{N}$ is a fixed natural number. Also define $N = n^2$. For simplicity, we assume that $n \geq 4$.

The proof is by induction. We begin with a small base graph with one interval which has its own dedicated path (that is, whenever the value of λ lies in that interval, its dedicated path is the shortest path in the graph). In each inductive step, we roughly triple the size of the graph and subdivide each interval into n intervals. Each of those intervals have their own dedicated paths. After $\log n$ steps, we end up with a graph on $\text{poly}(n)$ vertices with $n^{\log n}$ intervals.

A. Inductive Definition of Intervals

In our recursive construction, we will construct paths that reign as the shortest path in particular intervals (as mentioned above, each interval has its dedicated path) for the parameter λ . We will construct a large number of intervals and show that a different path is the unique shortest path in each interval. This will establish that there are many pieces in the piece-wise linear cost of the shortest path in our graph. The construction of the graph and the role of the intervals is described in detail later. We first place the framework by describing the intervals inductively. The intervals depend on the parameter N (recall $N = n^2$). Then, for $m = 0, 1, \dots, \lfloor \log n \rfloor$,

and $j = 0, 1, \dots, n^m - 1$, we define points $\alpha(j, m) \in \mathbb{Q}$ inductively; these points will be used to define the intervals. Each interval is of length $N - 2$. At the m -th step of our construction, we have n^m intervals.

Base case ($m = 0$): We set $\alpha(0, 0) = 0$. (Since $0 \leq j \leq n^m - 1$, the only possible value for j is 0.)

Induction ($m \geq 1$): For $m \geq 1$ and $0 \leq j \leq n^m - 1$, we write $j = nd + r$, where $0 \leq d \leq n^{m-1} - 1$ and $0 \leq r \leq n - 1$; then, we set $\alpha(j, m) = N\alpha(d, m - 1) + N(r + 1)$.

Intervals: For $m \geq 0$ and $0 \leq j \leq n^m - 1$, let $\mathcal{I}(j, m) = [\alpha(j, m) + 1, \alpha(j, m) + N - 1]$.

The idea behind this definition of the intervals is as follows. Given a set of n^{m-1} intervals at level $m - 1$, we first *stretch* them by a factor N (the corresponding graph can also be appropriately stretched by a factor N by replacing the λ in each edge weight with λ/N ; this will be explained in detail later). Then, we subdivide each interval into n disjoint intervals to obtain n^m intervals at level m .

However, in order to apply the induction hypothesis cleanly, we require that all of the n subdivided intervals are contained in the stretched version of their parent interval. Our definition of $\alpha(j, m)$ has $N(r + 1)$ instead of Nr precisely to ensure that $\mathcal{I}(j, m) \subseteq N \cdot \mathcal{I}(d, m - 1)$ (the definition in Mulmuley & Shah [3] unfortunately overlooks this point). Let us now summarize these observations.

- 1) For each $m \geq 0$ and $0 \leq j \leq n^m - 1$, $|\mathcal{I}(j, m)| = N - 2$.
- 2) For each $m \geq 1$ and $0 \leq j \leq n^m - 1$, $\mathcal{I}(j, m) \subseteq N \cdot \mathcal{I}(d, m - 1)$, where $d = \lfloor j/n \rfloor$.
- 3) For each $m \geq 0$, the intervals in the list $(\mathcal{I}(j, m) : j = 0, 1, \dots, n^m - 1)$ are disjoint.
- 4) For each $m \geq 0$ and $0 \leq j \leq n^m - 1$, $\mathcal{I}(j, m) \subseteq [0, N^{m+1}]$. In particular, $\alpha(j, m) \leq N^{m+1}$.

B. Inductive Construction of Graphs

Our induction depends on three parameters B , D and m , which impose constraints on the layered, weighted, planar graphs we construct.

The parameter B : $B \in \mathbb{N}$ denotes the number of vertices in the first (input) layer of this graph. B takes values of the form $1, n + 1, 2n + 1, \dots$, and for a fixed B , the variable $b \in \{0, 1, \dots, B - 1\}$ denotes an input vertex. All our paths originate in the first layer of the graph and end in the last layer. When we derive our main theorem from this construction, we set $B = 1$, which means our final graph has *one* input vertex. We call this unique input vertex s and connect all vertices in the last layer to a new vertex t using edges of weight 0, so that we have pristine s - t paths as promised. Thus, we do not mention s and t for the rest of our proof.

The parameter D : $D \in \mathbb{Q}$, $|D| \leq 1$ is used to determine the weights of the edges.

The parameter m : $m \in \mathbb{N}$, $m \geq 0$ is the induction parameter (this is the same m which is used to define the intervals).

A step-by-step visualization of this construction can be found in [23, Figure 3]. The formal induction will be carried out using a predicate Φ , which we now define.

The Predicate Φ

For B , D , m as described above, we say that *the predicate $\Phi(B, D, m)$ holds* if there is a layered, weighted, planar graph $G(B, D, m)$ with B input vertices, rational edge weights, and Bn^m paths P_{bj} (for $b = 0, 1, \dots, B - 1$ and $j = 0, 1, \dots, n^m - 1$) satisfying the following properties.

- (i) The graph $G(B, D, m)$ has at most $(3^{m+1} - 1)(B + mn)^4$ vertices.
- (ii) The weight of each edge e in the graph $G(B, D, m)$ has the form $a_e + b_e \lambda$ such that

$$a_e = a_{1e} + a_{2e}D \quad \text{and} \quad b_e = b_{1e} + b_{2e}D,$$

where $a_{1e}, b_{1e}, a_{2e}, b_{2e}$ are rational numbers with denominator at most $2^m n^2$ and numerator at most $(400NB)^{5m^2}$, in absolute value.

- (iii) For all b , j , and $\lambda \in \mathcal{I}(j, m)$, the *unique* shortest path from the input vertex b to the last layer of G is P_{bj} , and $\mathcal{C}(Q_b)(\lambda) - \mathcal{C}(P_{bj})(\lambda) \geq 1$, for all other paths Q_b from the input vertex b to the last layer of G .
- (iv) For all b and j ,

$$\mathcal{C}(P_{bj})(\lambda) = \mathcal{C}(P_{0j})(\lambda) + bD\alpha(j, m).$$

- (v) For all j , the paths $(P_{bj} : b = 0, 1, \dots, B - 1)$ are vertex-disjoint.
- (vi) For all b , the paths $(P_{bj} : j = 0, 1, \dots, n^m - 1)$ are distinct.

The following lemma is essentially the same as Lemma 4.1 of Mulmuley & Shah [3]. We closely follow their argument, slightly simplifying the induction, providing more detailed calculations, and correcting some errors; we crucially employ the planarized linking gadget of Section II and Lemma 9 to ensure that our graphs are planar.

Lemma 11 (Main lemma). *For all integers $B \geq 1$, rational numbers $D \in [-1, +1]$ and integers $m \geq 0$, the predicate $\Phi(B, D, m)$ holds.*

We will prove this lemma after using it to establish our main theorem.

Proof of Theorem 1. Taking $B = 1$, $D = 0$ and $m = \lfloor \log n \rfloor$, we conclude that $\Phi(1, 0, \lfloor \log n \rfloor)$ holds (Lemma 11). Using property (i), the number of vertices in the corresponding graph $G(1, 0, \lfloor \log n \rfloor)$ is at most

$$\begin{aligned} (3^{m+1} - 1)(B + mn)^4 &\leq (3^{\log n + 1} - 1)(1 + (\log n)n)^4 \\ &\leq (3n^{1.585})(1 + (\log n)n)^4 \\ &\leq 6n^{1.585}(n \log n)^4 \quad (\text{since } n \geq 4) \\ &\leq 6n^{1.585}(n \cdot n^{0.6})^4 \\ &\leq 6n^8. \end{aligned}$$

To this graph we attach a sink vertex t as stated earlier. The graph admits n^m disjoint intervals, with a different unique shortest s - t path in each (properties (iii),(vi)); so the cost of the shortest s - t path in this graph has $n^{\lfloor \log n \rfloor}$ pieces in the lower envelope.

Using property (ii) and substituting $B = 1$, $D = 0$ and $m = \lfloor \log n \rfloor$, the value of the largest coefficient (numerator or denominator) in the edge weights of the graph is at most

$$\begin{aligned} (400NB)^{5m^2} &\leq (400n^2)^{5(\log n)^2} \\ &\leq (400 \cdot 2^{2\log n})^{5(\log n)^2} \\ &\leq (2^{5\log n} \cdot 2^{2\log n})^{5(\log n)^2} \quad (\text{since } n \geq 4) \\ &\leq 2^{35(\log n)^3}. \end{aligned}$$

This implies that the bit lengths of the coefficients in the edge weights are bounded by $35(\log n)^3$.

Let ν be a large positive integer. Let n be the largest integer such that $6n^8 + 1 \leq \nu$. Note $n = \nu^{\Theta(1)}$ and hence $\log n = \Theta(\log \nu)$. Using the construction above (adding dummy isolated vertices if necessary), we obtain a graph on ν vertices, whose edge weights have rational coefficients with numerator and denominator of bit lengths bounded by $O((\log \nu)^3)$, and in which the cost of the shortest path has $\nu^{\Omega(\log \nu)}$ pieces in the lower envelope. This completes the proof of our main theorem.

Remark: Since we require integer edge weights, we can consider clearing all denominators in the coefficients. However, the LCM of the denominators may be prohibitively large. To keep the numbers small, we can modify our construction slightly. The points in the planar linking gadgets can be located at nearby points whose coordinates are multiples of (say) n^{-4} . This will ensure that the LCM of the denominators can be written in $O(\log \nu)$ bits. Clearing the denominators now keeps the final integer coefficients $O((\log \nu)^3)$ bits long.

Note that we did not use property (iv) or (v), which are needed merely in the inductive proof of the main lemma. \square

C. Proof of the Main Lemma

Proof of Lemma 11. We will use induction on m to verify $\Phi(B, D, m)$. For the base case ($m = 0$), let G consist of B disjoint edges, each of weight 0, leaving the B input vertices. The only choice for j in this case is $j = 0$ (since j varies from 0 to $n^m - 1$). To verify property (ii), note that each edge weight is of the form $((0 + 0 \cdot D) + (0 + 0 \cdot D)\lambda)$. To verify property (iv), recall that $\alpha(0, 0) = 0$. All the other properties for Φ are also easily verified.

Let $m \geq 1$. Assume that $\Phi(B', D', m - 1)$ holds for all B' and D' . We now fix B and D and show that $\Phi(B, D, m)$ holds for the graph $G(B, D, m)$. Based on B , D and m , we fix constants

$$K_L = 400N^{m+4}B^2; \quad (12)$$

$$K_R = 20N^3B; \quad (13)$$

$$D_L = \frac{N}{2K_L} \left(D - \frac{K_R}{N} \right); \quad (14)$$

$$D_R = 1. \quad (15)$$

These constants, which may seem mysterious, will be justified by the claims that follow. Let us now explain the construction and edge weights of G .

Construction of G : The graph G is built by serially concatenating four components: G^L , G^M , LINK and G^R . Let G^L be the graph corresponding to the induction hypothesis $\Phi(B, D_L, m - 1)$; we refer to the corresponding Bn^{m-1} paths by P_{bd}^L where $0 \leq b \leq B - 1$ and $0 \leq d \leq n^{m-1} - 1$. Let G^M denote the graph obtained by mirroring G^L about its last layer and reversing the directions of its edges so that all edges go from left to right (see Figure 3). Thus, G^M has B vertices in its last layer. Let G^R be the graph corresponding to the induction hypothesis $\Phi(B + n, D_R, m - 1)$; we refer to the corresponding $(B + n)n^{m-1}$ paths by P_{bd}^R where $0 \leq b \leq B + n - 1$ and $0 \leq d \leq n^{m-1} - 1$.

Edge weights of G : We need to transform the edges weights in G^L , G^M and G^R before we put them together with a linking gadget to obtain our graph G . Let w_e^L , w_e^M and w_e^R denote the weights of the edges in G^L , G^M and G^R , and let w_e denote their weights in G .

$$\begin{aligned} w_e(\lambda) &\leftarrow K_L \cdot w_e^L(\lambda/N) & \forall e \in E(G^L) \\ w_e(\lambda) &\leftarrow K_L \cdot w_e^M(\lambda/N) & \forall e \in E(G^M) \\ w_e(\lambda) &\leftarrow K_R \cdot w_e^R(\lambda/N) & \forall e \in E(G^R) \end{aligned}$$

In essence, we are *scaling* (by factors K_L and K_R) and *stretching* (by a factor N)² our already existing solutions for G^L , G^M and G^R so that together they can form a solution for G .

Linking gadget: Let $\mathbf{L}(B, n)$ be the non-planar linking gadget with edge weights

$$w_{b,b+r} = NDrb + \frac{K_R}{N} \left(\left(\frac{r(r+1)}{2} \right) N - r\lambda \right), \quad (16)$$

where $0 \leq b \leq B - 1$ and $0 \leq r \leq n$. Let $\mathbf{L}^{\text{pl}}(B, n)$ be its planarized version. Note that the D used in (16) is the D that was part of the predicate $\Phi(B, D, m)$ (neither D_L nor D_R). The graph G obtained by composing G^L , G^M , \mathbf{L}^{pl} and G^R is shown in Figure 3. Since G^L , G^M and G^R are planar by induction, and $\mathbf{L}^{\text{pl}}(B, n)$ is planar, the graph obtained by composing them is also planar. This completes the description of all the constituent components of G .

Before we proceed further, let us verify that for our choice of parameters, \mathbf{L}^{pl} faithfully simulates its non-planar counterpart. Invoke Lemma 9 with $J(b, b+r) = NDrb$, $K = K_R$ and $L = -K_R/N$. For the setting of K_R in (13), we have the following (recall $N = n^2$).

$$\begin{aligned} n^2 \left(1 + 2 \max_{e \in E(\mathbf{L}^{\text{pl}}(B, n))} |J(e)| \right) &\leq n^2 (1 + 2N|D|nB) \\ &\leq 4(|D| + 1)N^{2.5}B \\ &\leq K_R. \end{aligned}$$

Thus, the requirement $K \geq n^2 (1 + 2 \max_e |J(e)|)$ of Lemma 9 holds. Therefore, $\mathbf{L}^{\text{pl}}(B, n)$ faithfully simulates $\mathbf{L}(B, n)$.

²Recall that our intervals are stretched by a factor N when we go from one level of recursion to the next (Subsection III-A).

Finally, in order to invoke the induction hypothesis for $\Phi(B, D_L, m-1)$ and $\Phi(B+n, D_R, m-1)$, we need to show that $|D_L| \leq 1$ ($|D_R| = 1$ from (15)).

$$\begin{aligned} |D_L| &= \left| \frac{N}{2K_L}(D - 20N^2B) \right| && \text{(from (14))} \\ &\leq \left| \frac{ND}{800N^{m+4}B^2} \right| + \left| \frac{20N^3B}{800N^{m+4}B^2} \right| && \text{(from (12))} \\ &\leq \frac{1}{800} + \frac{1}{40} \ll 1. && \text{(since } |D| \leq 1) \end{aligned}$$

Thus, we can work under the assumption that $\Phi(B, D_L, m-1)$ and $\Phi(B+n, D_R, m-1)$ hold. We may view G as (see Figure 3)

$$G = G^L \circ G^M \circ \mathbf{L}^{\text{pl}} \circ G^R,$$

where \circ represents concatenation of graphs. To show that $\Phi(B, D, m)$ holds, we will first show through calculations that properties (i), (ii) hold. To verify that properties (iii), (iv), (v), (vi) hold, we will exhibit Bn^m paths in G . For $0 \leq j \leq n^m - 1$, write $j = nd + r$ with $0 \leq d \leq n^{m-1} - 1$ and $0 \leq r \leq n - 1$; then for $0 \leq b \leq B - 1$, let

$$P_{bj} = P_{bd}^L \circ (P_{bd}^L)^{\text{rev}} \circ \text{link}(b, b+r+1) \circ P_{b+r+1, d}^R, \quad (17)$$

where $\text{link}(b, b+r+1)$ is the unique shortest path (the straight line) in \mathbf{L}^{pl} connecting vertex b in the last layer of G^M to vertex $b+r+1$ in the first layer of G^R . We will show that in the interval $\mathcal{I}(j, m)$, the path P_{bj} as defined by (17) is the shortest path from the input vertex b in G . We are now set to show that properties (i) through (vi) hold for Φ .

Property (i): Note that the planarization of the linking gadget $\mathbf{L}^{\text{pl}}(B, n)$ adds at most $(B+n)^4$ new vertices. This means that the number of vertices in the planarized version of G is at most

$$\begin{aligned} &\underbrace{2(3^m - 1)(B + (m-1)n)^4}_{G^L, G^M} + \underbrace{(B+n)^4}_{\mathbf{L}^{\text{pl}}(B, n)} \\ &+ \underbrace{(3^m - 1)(B + n + (m-1)n)^4}_{G^R}, \end{aligned}$$

which is at most $(3^{m+1} - 1)(B + mn)^4$. Thus, property (i) holds. We now verify property (ii).

Property (ii): Using the induction hypothesis, we know that each edge e in the graph $G(B', D', m-1)$ has the form $a_e + b_e\lambda$, where $a_e = a_{1e} + a_{2e}D$, $b_e = b_{1e} + b_{2e}D$. Also,

$$\begin{aligned} \max_e \{ |\text{num}(a_{1e})|, |\text{num}(a_{2e})|, |\text{num}(b_{1e})|, |\text{num}(b_{2e})| \} &\leq (400NB')^{5(m-1)^2}, \\ \max_e \{ |\text{den}(a_{1e})|, |\text{den}(a_{2e})|, |\text{den}(b_{1e})|, |\text{den}(b_{2e})| \} &\leq 2^{m-1}n^2, \end{aligned}$$

where e ranges over all the edges of $G(B', D', m-1)$, num stands for numerator, and den stands for denominator. Each edge of G comes from either G^L , G^M , $\mathbf{L}^{\text{pl}}(B, n)$ or G^R .

First we consider edges coming from G^L (we do not consider G^M separately it has the same edge weights as G^L).

Let e be an edge of G coming from G^L . Using the induction hypothesis,

$$\begin{aligned} a_e^L &= a_{1e}^L + a_{2e}^L \left(\frac{N}{2K_L} \left(D - \frac{K_R}{N} \right) \right). && \text{(before scaling)} \\ b_e^L &= b_{1e}^L + b_{2e}^L \left(\frac{N}{2K_L} \left(D - \frac{K_R}{N} \right) \right). && \text{(before scaling)} \end{aligned}$$

However, once G^L becomes a part of G , a_e^L is scaled by K_L and b_e^L is scaled by K_L/N .

$$\begin{aligned} a_e &= K_L a_{1e}^L + \frac{a_{2e}^L N}{2} \left(D - \frac{K_R}{N} \right) && \text{(after scaling)} \\ &= 400 N^{m+4} B^2 a_{1e}^L + \frac{a_{2e}^L N}{2} \left(D - \frac{20 N^3 B}{N} \right) && \text{(from (12))} \\ &= \underbrace{(400 N^{m+4} B^2 a_{1e}^L - 10 a_{2e}^L N^3 B)}_{a_{1e}} + \underbrace{\left(\frac{a_{2e}^L N}{2} \right)}_{a_{2e}} D. \\ b_e &= \frac{K_L b_{1e}^L}{N} + \frac{b_{2e}^L}{2} \left(D - \frac{K_R}{N} \right) && \text{(after scaling)} \\ &= \frac{400 N^{m+4} B^2 b_{1e}^L}{N} + \frac{b_{2e}^L}{2} \left(D - \frac{20 N^3 B}{N} \right) && \text{(from (12))} \\ &= \underbrace{(400 N^{m+3} B^2 b_{1e}^L - 10 b_{2e}^L N^2 B)}_{b_{1e}} + \underbrace{\left(\frac{b_{2e}^L}{2} \right)}_{b_{2e}} D. \end{aligned}$$

Thus, a_e and b_e have the required form. If the denominators of $a_{1e}^L, a_{2e}^L, b_{1e}^L$ and b_{2e}^L have absolute value at most $2^{m-1}n^2$, then the denominators of a_{1e}, a_{2e}, b_{1e} and b_{2e} have absolute value at most $2^m n^2$. Now we need to check for the numerators.

$$\begin{aligned} |\text{num}(a_{1e})| &= |\text{num}(400 N^{m+4} B^2 a_{1e}^L - 10 a_{2e}^L N^3 B)| \\ &\leq \left| 400 N^{m+4} B^2 (400NB)^{5(m-1)^2} \right| \\ &\quad + \left| 10 (400NB)^{5(m-1)^2} N^3 B \right| \\ &\leq \left| (200NB)^{10m-5} (400NB)^{5(m-1)^2} \right| \\ &\quad + \left| (200NB)^{10m-5} (400NB)^{5(m-1)^2} \right| \\ &\leq (400NB)^{5m^2}. \end{aligned}$$

We skip the proof for a_{2e}, b_{1e} and b_{2e} . Now we consider edges coming from $\mathbf{L}^{\text{pl}}(B, n)$. Let $(b, b+r) \in E(\mathbf{L}(B, n))$.

$$\begin{aligned} w_{b, b+r} &= NDrb + \frac{K_R}{N} \left(\left(\frac{r(r+1)}{2} \right) N - r\lambda \right) && \text{(from (16))} \\ &= NDrb + \frac{r(r+1)K_R}{2} - \frac{rK_R}{N}\lambda \\ &= \underbrace{(NDrb + 10r(r+1)N^3B)}_{a_e} + \underbrace{(-20rN^2B)\lambda}_{b_e}. \end{aligned}$$

Note that all these coefficients are integers. However, these are the edge weights from the *non-planar* linking gadget. Once we planarize it, the edges in the *planar* linking gadget can have

denominators at most n^2 (see Fact 7). As for the numerator (recall that $0 \leq b \leq B-1$, $0 \leq r \leq n$ and $N = n^2$),

$$\begin{aligned} |NDrb + 10r(r+1)N^3B| &\leq NDnB + 10N^4B \\ &\leq (400NB)^{5m^2}. \end{aligned}$$

Finally we consider edges coming from G^R . Let e be an edge of G coming from G^R . Using the induction hypothesis, we have the following.

$$\begin{aligned} a_e^R &= a_{1e}^R + a_{2e}^R D_R. && \text{(before scaling)} \\ b_e^R &= b_{1e}^R + b_{2e}^R D_R. && \text{(before scaling)} \\ a_e &= K_R a_{1e}^R + K_R a_{2e}^R && \text{(after scaling)} \\ &= \underbrace{20N^3 B a_{1e}^R + 20N^3 B a_{2e}^R}_{a_{1e}}. && \text{(using (13))} \\ b_e &= \frac{K_R}{N} b_{1e}^R + \frac{K_R}{N} b_{2e}^R && \text{(after scaling)} \\ &= \underbrace{20N^2 B a_{1e}^R + 20N^2 B a_{2e}^R}_{b_{1e}}. && \text{(using (13))} \end{aligned}$$

Thus, $a_e = a_{1e} + 0 \cdot D$ and $b_e = b_{1e} + 0 \cdot D$ have the required form. Also, a_{1e} and b_{1e} are integers.

$$\begin{aligned} |\text{num}(a_{1e})| &= |20N^3 B a_{1e}^R + 20N^3 B a_{2e}^R| \\ &\leq |20N^3 B (400NB)^{5(m-1)^2}| \\ &\quad + |20N^3 B (400NB)^{5(m-1)^2}| \\ &\leq (400NB)^{5m^2}. \end{aligned}$$

We skip the proof for b_{1e} . This finishes the verification of property (ii).

Properties (v), (vi): Given our definition of P_{bj} (17), properties (v) and (vi) are straightforward to verify.

Property (iv): In the subsequent calculations, we use the following notation. Paths of G are composed of paths from G^L , G^M and G^R ; we use \mathcal{C}_L , \mathcal{C}_M and \mathcal{C}_R to denote the costs of those subpaths in their constituent graphs. For example, $\mathcal{C}_L(P_{bd}^L)(\lambda)$ denotes the cost of the path P_{bd}^L as a function of λ in the graph G^L . When G^L is used as a component in G , this cost is scaled by a factor of K_L and stretched by a factor N . Thus, the cost of the path P_{bj} in the graph G is given by

$$\begin{aligned} \mathcal{C}(P_{bj})(\lambda) &= K_L \mathcal{C}_L(P_{bd}^L) \left(\frac{\lambda}{N} \right) + K_L \mathcal{C}_M((P_{bd}^L)^{\text{rev}}) \left(\frac{\lambda}{N} \right) \\ &\quad + w_{b,b+r+1} + K_R \mathcal{C}_R(P_{b+r+1,d}^R) \left(\frac{\lambda}{N} \right) \\ &= 2K_L \mathcal{C}_L(P_{bd}^L) \left(\frac{\lambda}{N} \right) \\ &\quad + w_{b,b+r+1} + K_R \mathcal{C}_R(P_{b+r+1,d}^R) \left(\frac{\lambda}{N} \right) \\ &= 2K_L \left[\mathcal{C}_L(P_{0d}^L) \left(\frac{\lambda}{N} \right) + b D_L \alpha(d, m-1) \right] \\ &\quad + ND(r+1)b + \frac{K_R}{N} \left[\frac{(r+1)(r+2)}{2} N - (r+1)\lambda \right] \\ &\quad + K_R \left[\mathcal{C}_R(P_{0d}^R) \left(\frac{\lambda}{N} \right) + (b+r+1) D_R \alpha(d, m-1) \right]. \end{aligned}$$

Substitute $b = 0$ to get

$$\begin{aligned} \mathcal{C}(P_{0j})(\lambda) &= 2K_L \mathcal{C}_L(P_{0d}^L) \left(\frac{\lambda}{N} \right) \\ &\quad + \frac{K_R}{N} \left[\frac{(r+1)(r+2)}{2} N - (r+1)\lambda \right] \\ &\quad + K_R \mathcal{C}_R(P_{0d}^R) \left(\frac{\lambda}{N} \right) + K_R(r+1) D_R \alpha(d, m-1). \end{aligned}$$

With this expression for $\mathcal{C}(P_{0j})(\lambda)$, we obtain

$$\begin{aligned} \mathcal{C}(P_{bj})(\lambda) &= \mathcal{C}(P_{0j})(\lambda) + b[2K_L D_L \alpha(d, m-1) \\ &\quad + K_R D_R \alpha(d, m-1) + ND(r+1)] \\ &= \mathcal{C}(P_{0j})(\lambda) + b \left[2K_L \frac{N}{2K_L} \left(D - \frac{K_R}{N} \right) \alpha(d, m-1) \right. \\ &\quad \left. + K_R \alpha(d, m-1) + ND(r+1) \right] \\ &= \mathcal{C}(P_{0j})(\lambda) + b[ND\alpha(d, m-1) - K_R \alpha(d, m-1) \\ &\quad + K_R \alpha(d, m-1) + ND(r+1)] \\ &= \mathcal{C}(P_{0j})(\lambda) + bD[N\alpha(d, m-1) + N(r+1)] \\ &= \mathcal{C}(P_{0j})(\lambda) + bD\alpha(j, m). \end{aligned}$$

Thus, property (iv) also holds. All that remains is to verify property (iii).

Property (iii): To verify property (iii), we need to check that P_{bj} as defined above is indeed the shortest path from input vertex b to the last layer when $\lambda \in \mathcal{I}(j, m)$, and any deviation from it attracts significant additional cost. We do this through two claims (Claim 18 and Claim 19).

In Claim 18, we track paths from an input vertex as they travel through G^L and G^M . In Claim 19, we analyze how such paths continue through \mathbf{L}^{Pl} and G^R . Fix a j ($0 \leq j \leq n^m - 1$), say $j = nd + r$, for $0 \leq d \leq n^{m-1} - 1$ and $0 \leq r \leq n - 1$ and a $\lambda \in \mathcal{I}(j, m)$. Note that since $\lambda \in \mathcal{I}(j, m)$, we have $\lambda/N \in \mathcal{I}(d, m-1) = [\alpha(d, m-1) + 1, \alpha(d, m-1) + N - 1]$.

Claim 18. *Let Q be a path from the input vertex b to the last layer of $G^L \circ G^M$ (note that $P_{bd}^L \circ (P_{bd}^L)^{\text{rev}}$ is one such path). If $Q \neq P_{bd}^L \circ (P_{bd}^L)^{\text{rev}}$, then*

$$\mathcal{C}(Q)(\lambda) - \mathcal{C}(P_{bd}^L \circ (P_{bd}^L)^{\text{rev}})(\lambda) \geq K_L/2.$$

Proof of claim. We omit the argument λ in this discussion. Let $Q = Q^L \circ Q^M$, where Q^L is the subpath of Q in G^L and Q^M is the subpath of Q in G^M . Suppose Q^M terminates at vertex c in the last layer of G^M . Then,

$$\begin{aligned} \mathcal{C}(Q) - \mathcal{C}(P_{bd}^L \circ (P_{bd}^L)^{\text{rev}}) &= (\mathcal{C}(Q^L) + \mathcal{C}(Q^M)) - (\mathcal{C}(P_{bd}^L) + \mathcal{C}((P_{bd}^L)^{\text{rev}})) \\ &= (\mathcal{C}(Q^L) - \mathcal{C}(P_{bd}^L)) + (\mathcal{C}(Q^M) - \mathcal{C}(P_{cd}^L)) \\ &\quad + (\mathcal{C}((P_{cd}^L)^{\text{rev}}) - \mathcal{C}((P_{bd}^L)^{\text{rev}})) \\ &\geq \underbrace{\mathcal{C}(Q^L) - \mathcal{C}(P_{bd}^L)}_{\text{Term I}} + \underbrace{\mathcal{C}(Q^M) - \mathcal{C}(P_{cd}^L)}_{\text{Term II}} \\ &\quad - \underbrace{K_L B D_L \alpha(d, m-1)}_{\text{Term III}}. \end{aligned}$$

To obtain Term III, we use part (ii) of the induction hypothesis for G^L , whose edge costs we evaluated at λ/N and scaled by K_L ; recall that $\lambda/N \in \mathcal{I}(d, m-1)$. If $Q \neq P_{bd}^L \circ (P_{bd}^L)^{\text{rev}}$, then one of the following is true.

- (a) $Q^L \neq P_{bd}^L$;
- (b) $c = b$ and $Q^M \neq (P_{bd}^L)^{\text{rev}}$;
- (c) $c \neq b$ and $Q^M \neq (P_{cd}^L)^{\text{rev}}$ (here we use the fact that the paths P_{bd}^L and P_{cd}^L are vertex-disjoint if $c \neq b$).

From property (iii) of the induction hypothesis, the costs of a shortest and a non-shortest path from the same input vertex differ by at least one in G^L and G^M ; after scaling all the edge weights of G^L and G^M by a factor of K_L , this difference becomes at least K_L . Also note that both Term I and Term II are non-negative. Thus we can conclude the following.

If (a) is true, then Term I $\geq K_L$. If (b) or (c) is true, then Term II $\geq K_L$. Recall that $\alpha(d, m-1) \leq N^m$. For the setting of K_L according to (12), we have $|\text{Term III}| = |K_L B D_L \alpha(d, m-1)| \ll K_L/10$. This completes the proof of Claim 18. \square

Since K_L is positive, Claim 18 implies that $P_{bd}^L \circ (P_{bd}^L)^{\text{rev}}$ is the shortest path from the input vertex b to the last layer of $G^L \circ G^M$. Now, we need to argue that the overall shortest path must be an extension of this. The next claim shows that the shortest path from an input vertex b of $\mathbf{L}(B, n)$ (note that is the non-planar version of the linking gadget) in the graph $\mathbf{L}(B, n) \circ G^R$ must follow the route prescribed by (17).

Claim 19. *Let $\lambda \in \mathcal{I}(j, m)$, where $j = nd + r$ ($0 \leq d \leq n^{m-1} - 1$ and $0 \leq r \leq n - 1$). Let P be a path from the input vertex b of $\mathbf{L}(B, n)$ to the last layer of G^R (note that $\text{link}(b, b+r+1) \circ P_{b+r+1, d}^R$ is one such path). If $P \neq \text{link}(b, b+r+1) \circ P_{b+r+1, d}^R$ then*

$$\mathcal{C}(P)(\lambda) - \mathcal{C}(\text{link}(b, b+r+1) \circ P_{b+r+1, d}^R)(\lambda) \geq 1.$$

Proof of claim. Fix the input vertex b . The induction hypothesis guarantees that $P_{x, d}^R$ is the unique shortest path from the input vertex x of G^R to the last layer of G^R . We may assume that P travels along the shortest path in G^R , that is, it has the form

$$P_k = \text{link}(b, b+k) \circ P_{b+k, d}^R,$$

for some $k \in \{0, 1, \dots, n\}$. Let $Z_k = \mathcal{C}(P_k)(\lambda)$. We will show that for $\lambda \in \mathcal{I}(j, m)$, we have

$$Z_0 \gg Z_1 \gg \dots \gg Z_r \gg Z_{r+1} \ll Z_{r+2} \ll \dots \ll Z_n, \quad (20)$$

where we use \gg and \ll to suggest that there is a large gaps between the quantities. Our proof strategy is to compare successive values of Z_k . We will show that $Z_k - Z_{k-1}$ is negative whenever $k \leq r+1$ and positive otherwise. Indeed, for $k = 1, 2, \dots, n$, we have

$$Z_k - Z_{k-1} = w_{b, b+k} - w_{b, b+k-1} + \mathcal{C}(P_{b+k, d}^R)(\lambda) - \mathcal{C}(P_{b+k-1, d}^R)(\lambda),$$

where

$$w_{b, b+k} = NDkb + \frac{K_R}{N} \left(\binom{k(k+1)}{2} N - k\lambda \right);$$

$$\mathcal{C}(P_{b+k, d}^R)(\lambda) = K_R [\mathcal{C}_R(P_{0, d}^R)(\lambda/N) + (b+k)D_R\alpha(d, m-1)].$$

$w_{b, b+k-1}$ and $\mathcal{C}(P_{b+k-1, d}^R)(\lambda)$ are defined similarly. Thus,

$$Z_k - Z_{k-1} = NDb + \frac{K_R}{N}(kN - \lambda) + K_R D_R \alpha(d, m-1) \quad (21)$$

$$= NDb + \frac{K_R}{N}(kN + N\alpha(d, m-1) - \lambda) \quad (22)$$

$$= NDb + \frac{K_R}{N}(\alpha(k-1, m) - \lambda). \quad (23)$$

Since $\lambda \in \mathcal{I}(r, m) = [\alpha(r, m) + 1, \alpha(r, m) + N - 1]$, we have

$$\begin{aligned} \alpha(k-1, m) - \lambda &\in [\alpha(k-1, m) - \alpha(r, m) - N + 1, \\ &\alpha(k-1, m) - \alpha(r, m) - 1] \\ &= [(k - (r+1))N - N + 1, \\ &(k - (r+1))N - 1]. \end{aligned}$$

Thus, for $k = 1, 2, \dots, r+1$, we have $\alpha(k-1, m) - \lambda \leq -1$ and for $k = r+2, \dots, n$, we have $\alpha(k-1, m) - \lambda \geq +1$. Returning to (23) with this, we obtain

$$Z_k - Z_{k-1} \leq NDb - \frac{K_R}{N} \quad \text{for } k = 1, \dots, r+1, \text{ and} \quad (24)$$

$$Z_k - Z_{k-1} \geq NDb + \frac{K_R}{N} \quad \text{for } k = r+2, \dots, n. \quad (25)$$

Since $K_R \gg N^2b$ and $-1 \leq D \leq +1$, the RHS of (24) is negative and the RHS of (25) is positive. This confirms (20) and establishes Claim 19. \square

We are now in a position to establish property (iii) and complete the induction. By Claim 18, if the subpath of the shortest path from b in $G^L \circ G^M$ is not $P_{bd}^L \circ (P_{bd}^L)^{\text{rev}}$, then the increase in cost is at least $K_L/2$. We show that the difference in cost between every two paths originating at an input vertex of $\mathbf{L}(B, n)$ and terminating in the last layer of G^R is much smaller than $K_L/2$; this forces the subpath of the shortest path from b in $G^L \circ G^M$ to be precisely $P_{bd}^L \circ (P_{bd}^L)^{\text{rev}}$. Let P_1 and P_2 be paths originating at an input vertex of $\mathbf{L}(B, n)$ and terminating in the last layer of G^R . Then,

$$\begin{aligned} &|\mathcal{C}(P_1)(\lambda) - \mathcal{C}(P_2)(\lambda)| \\ &\leq |NDnB| + \left| \frac{K_R}{N} (n^2N + n\lambda) \right| + |K_R D_B \alpha(d, m-1)| \\ &\leq |N^2DB| + \left| \frac{K_R}{N} (N^2 + n(\alpha(j, m) + N)) \right| \\ &\quad + |K_R D_B \alpha(d, m-1)| \\ &\leq |N^2DB| + \left| \frac{K_R}{N} (N^2 + n(N^{m+1} + N)) \right| + |K_R D_B N^m| \\ &\leq |N^2DB| + |K_R N^{m+1}| + |K_R D_B N^m| \\ &\ll K_L/10. \end{aligned}$$

Thus, the shortest path in G must follow the path $P_{bd}^L \circ (P_{bd}^R)^{\text{rev}}$ until it arrives at the first layer of $\mathbf{L}(B, n)$; for if it does not, then its cost is at least $K_L/2 - K_L/10 \gg 1$ more than the cost of P_{bd} . Claim 19 now confirms that it must continue by taking the edge $\text{link}(b, b+r+1)$ and $P_{b+r+1, d}^R$; any deviation from this path will incur an increase in cost of at least 1. Therefore, the shortest path in G in the interval $\mathcal{I}(j, m)$ is P_{bj} as promised (17). This completes the verification of property (iii), hence completing the proof of Lemma 11. \square

IV. THIN GRIDS

In this section, we will see a simple yet instructive application of alternation-free sequences (Proposition 5). In the forthcoming sections, we upper bound the parametric shortest path complexity of graphs under various settings. All our upper bound proofs, including this one, are based around the idea of alternation-free sequences. We later also present some lower bounds for alternation-free sequences (Section VII).

Definition 26. *The $p \times q$ directed grid graph, denoted by $\Upsilon_{p,q}$, is defined as follows.*

- (a) $V(\Upsilon_{p,q}) = \{(i, j) : 1 \leq i \leq p, 1 \leq j \leq q\}$.
- (b) $((i_1, j_1), (i_2, j_2)) \in E(\Upsilon_{p,q})^3$ if and only if $(i_1 = i_2 \text{ and } j_2 = j_1 + 1)$ or $(j_1 = j_2 \text{ and } i_2 = i_1 + 1)$.

In other words, the vertices of $\Upsilon_{p,q}$ form a 2D lattice, and a vertex is connected to the vertex immediately to its right and the vertex immediately above it. \diamond

Let $\varphi^{\text{gr}}(p, q, \beta)$ be the parametric shortest path complexity of $\Upsilon_{p,q}$ where the bit lengths of the coefficients in the weights of the edges are bounded by β . The planar graphs that we construct as part of our main result (Theorem 1) can be remodeled into grid graphs at the expense of a small (polynomial factor) blow-up in size of the graph. Thus, $\varphi^{\text{gr}}(n, n, O((\log n)^3)) \geq n^{\Omega(\log n)}$. This settles the parametric complexity for *square* grids. We ask the same question for thin *rectangular* grids. These are the graphs $\Upsilon_{p,q}$ with $p \ll q$. Note that $\varphi^{\text{gr}}(1, n, \infty) \leq 1$ and $\varphi^{\text{gr}}(2, n, \infty) \leq n$ trivially. The problem becomes nontrivial for $3 \times n$ grids. We have the following result.

Theorem 27. $\varphi^{\text{gr}}(3, n, \infty) \leq 5n$.

Proof. Our proof is via an upper bound on the maximum length of an alternation-free sequence of paths in $\Upsilon_{3,n}$. Now $\Upsilon_{3,n}$ has 3 rows and n columns; let the vertices in its middle row be $\{v_1, v_2, \dots, v_n\}$, arranged in increasing order of their distance from s . Our proof strategy is as follows. Given an alternation-free sequence of paths \mathcal{P} in $\Upsilon_{3,n}$, we will assign one v_i to each path in \mathcal{P} (different paths may be assigned the same v_i). Then we will show that each v_i can be assigned to at most 5 paths in \mathcal{P} , thus proving an upper bound of $5n$ on the length of \mathcal{P} .

Since every path from s to t must pass through the middle row, an s - t path may be defined by the two vertices it uses to enter and leave the middle row. More formally, for $1 \leq$

³The ordering of (i_1, j_1) and (i_2, j_2) is important since this is a directed graph.

$i \leq j \leq n$, let $P(i, j)$ be the path from s to t in which v_i is the first vertex of the middle row that lies on $P(i, j)$ and v_j is the last vertex of the middle row that lies on $P(i, j)$. Using this notation, let the alternation-free sequence be $\mathcal{P} = (P(i_1, j_1), P(i_2, j_2), \dots, P(i_T, j_T))$. We will prove that $T \leq 5n$.

We now describe how we assign a middle row vertex to each path in \mathcal{P} . For this, we will compare the k -th path of \mathcal{P} with all earlier paths of \mathcal{P} as follows. For each $k \in \{1, 2, \dots, T\}$, consider the maximum r ($1 \leq r \leq k-1$) such that $[i_r, j_r] \cap [i_k, j_k] \neq \emptyset$. Three cases arise.

- (a) If no such r exists, then assign v_{i_k} to $P(i_k, j_k)$.
- (b) If $i_r \neq i_k$, then assign v_ℓ to $P(i_k, j_k)$, where $\ell = \max\{i_r, i_k\}$.
- (c) If $(i_r = i_k \text{ and } j_r \neq j_k)$, then assign v_ℓ to $P(i_k, j_k)$, where $\ell = \min\{j_r, j_k\}$.

First, note that these are the only possible cases. If case (a) is false (that is, an r does exist), then at least one out of cases (b) or (c) is true, since all the paths in \mathcal{P} are distinct.

The crucial observation now is that, in $P(i_k, j_k)$, either the vertex v_ℓ appears for the first time in \mathcal{P} (case (a))⁴, or the incoming edge to v_ℓ has changed since its *most recent* occurrence in \mathcal{P} (case (b)), or the outgoing edge from v_ℓ has changed since its *most recent* occurrence in \mathcal{P} (case (c)). Fix a middle row vertex v_m . Clearly, v_m can appear for the first time in \mathcal{P} at most once. Also, once v_m has appeared in \mathcal{P} , the incoming and outgoing edges of v_m in later paths of \mathcal{P} can each change at most two times (see Claim 28 below). Thus, v_m can be assigned to at most 5 different paths in \mathcal{P} . Summing over all choices of v_m , we get $|\mathcal{P}| = T \leq 5n$. \square

Claim 28. *Let v_m and \mathcal{P} be as defined in the proof of Theorem 27. Then the incoming and outgoing edges of v_m in \mathcal{P} can each change at most two times.*

Proof. We will show that the incoming edge to v_m in \mathcal{P} can change at most two times. Let $\text{pred}_r(v_m)$ be the predecessor of v_m on the path $P(i_r, j_r)$. Since v_m has in-degree 2 (for $m > 1$), $\text{pred}_r(v_m)$ is either v_{m-1} or x , for some vertex x in the first row of $\Upsilon_{3,n}$. Note that the edge (x, v_m) fixes the s - v_m subpath, and changing the incoming edge of v_m in a later path in \mathcal{P} amounts to abandoning that s - v_m subpath. Therefore, the edge (x, v_m) does not occur in any subsequent path in \mathcal{P} . Let us now make this argument formal.

Suppose there exist four paths $P(i_a, j_a), P(i_b, j_b), P(i_c, j_c)$ and $P(i_d, j_d)$ in \mathcal{P} with $a < b < c < d$ such that $\text{pred}_a(v_m) = \text{pred}_c(v_m) = v_{m-1}$ and $\text{pred}_b(v_m) = \text{pred}_d(v_m) = x$. This means that the incoming edge to v_m has changed *three times*. Since there is a unique path from s to x in $\Upsilon_{3,n}$, we have $P(i_b, j_b)[s, v_m] = P(i_d, j_d)[s, v_m] \neq P(i_c, j_c)[s, v_m]$, implying that \mathcal{P} is not alternation-free, which is a contradiction.

It can also be shown that the outgoing edge from v_m in \mathcal{P} can change at most two times. We skip the proof because it is along similar lines. \square

⁴That is, v_ℓ is not part of $P(i_r, j_r)$, for all $1 \leq r \leq k-1$.

It is not known if $\varphi^{\text{gr}}(4, n, \infty) \leq O(n)$. However, a simple induction on the grid size shows the following generalization of Theorem 27.

Theorem 29. *For $3 \leq p \leq q$, we have $\varphi^{\text{gr}}(p, q, \infty) \leq O(q(\log q)^{p-3})$.*

Proof. The proof is by induction on p . For the base case ($p = 3$), Theorem 27 implies that $\varphi^{\text{gr}}(3, q, \infty) \leq O(q)$. For the inductive case, fix a value of p (where $4 \leq p \leq q$), and assume that $\varphi^{\text{gr}}(p', q, \infty) \leq O(q(\log q)^{p'-3})$ for all $3 \leq p' < p$. Now $\Upsilon_{p,q}$ has p rows and q columns; let the vertices in its $\lfloor \frac{q}{2} \rfloor$ -th column be $\{u_1, u_2, \dots, u_p\}$, arranged in increasing order of their distance from s . Our proof strategy is as follows. Let \mathcal{P} be the longest alternation-free sequence of paths in $\Upsilon_{p,q}$. We will partition \mathcal{P} into p alternation-free subsequences⁵ $\mathcal{P}_1, \mathcal{P}_2, \dots, \mathcal{P}_p$, and provide an upper bound for each. The sum of these p upper bounds is clearly an upper bound on $|\mathcal{P}|$.

Let $\mathcal{P} = \mathcal{P}_1 \cup \mathcal{P}_2 \cup \dots \cup \mathcal{P}_p$, where the sequence of paths in each \mathcal{P}_i respects its original ordering in \mathcal{P} . The partitions are defined as follows. For each path $P \in \mathcal{P}$, we have $P \in \mathcal{P}_i$ if and only if u_i is the first vertex of the $\lfloor \frac{q}{2} \rfloor$ -th column that lies on P . For each \mathcal{P}_i , we have

$$|\mathcal{P}_i| \leq \varphi^{\text{gr}}\left(i, \left\lfloor \frac{q}{2} \right\rfloor, \infty\right) + \varphi^{\text{gr}}\left(p - i + 1, q - \left\lfloor \frac{q}{2} \right\rfloor + 1, \infty\right).$$

We are now ready to provide an upper bound for $\varphi^{\text{gr}}(p, q, \infty)$. By definition, $\varphi^{\text{gr}}(p, q, \infty) = |\mathcal{P}| = \sum_{i=1}^p |\mathcal{P}_i|$. Thus,

$$\begin{aligned} \sum_{i=1}^p |\mathcal{P}_i| &\leq \sum_{i=1}^p \varphi^{\text{gr}}\left(i, \left\lfloor \frac{q}{2} \right\rfloor, \infty\right) \\ &\quad + \sum_{i=1}^p \varphi^{\text{gr}}\left(p - i + 1, q - \left\lfloor \frac{q}{2} \right\rfloor + 1, \infty\right) \\ &\leq 2 \sum_{i=1}^p \varphi^{\text{gr}}\left(i, q - \left\lfloor \frac{q}{2} \right\rfloor + 1, \infty\right) \\ &\leq 2 \underbrace{\varphi^{\text{gr}}\left(p, q - \left\lfloor \frac{q}{2} \right\rfloor + 1, \infty\right)}_{\text{Term I}} \\ &\quad + 2 \underbrace{\sum_{i=1}^{p-1} \varphi^{\text{gr}}\left(i, q - \left\lfloor \frac{q}{2} \right\rfloor + 1, \infty\right)}_{\text{Term II}}. \end{aligned}$$

Term II can be solved by invoking the induction hypothesis and Term I becomes part of the recurrence.

$$\begin{aligned} \varphi^{\text{gr}}(p, q, \infty) &\leq 2 \varphi^{\text{gr}}(p, r, \infty) + 2 \sum_{i=1}^{p-1} \varphi^{\text{gr}}(i, r, \infty) \\ &\leq 2 \varphi^{\text{gr}}(p, r, \infty) + 2c_1 \sum_{i=1}^{p-1} r(\log r)^{i-3} \\ &\leq 2 \varphi^{\text{gr}}(p, r, \infty) + 2c_1 r (c_2 (\log r)^{p-4}) \\ &\leq 2 \varphi^{\text{gr}}(p, r, \infty) + O(q(\log q)^{p-4}), \end{aligned}$$

⁵A subsequence of an alternation-free sequence is also alternation-free.

where c_1, c_2 are constants. Since r is roughly $q/2$, evaluating this final recurrence gives $\varphi^{\text{gr}}(p, q, \infty) \leq O(q(\log q)^{p-3})$. \square

Remark: Theorem 29 only helps for small values of p , that is, when $p \leq \frac{(\log q)^2}{\log \log q}$. For large p , the generalized upper bound of Gusfield gives a much better upper bound.

V. GRAPHS WITH EDGE WEIGHTS AS POLYNOMIALS

In this section, we consider graphs with edge weights of the form

$$w_e(\lambda) = a_d \lambda^d + a_{d-1} \lambda^{d-1} + \dots + a_1 \lambda + a_0.$$

The main result of this section is the following.

Theorem 2 (Upper bound with univariate polynomial edge weights).

$$\varphi^{\text{deg}(d)}(n) = n^{\log n + (\alpha(n) + O(1))^d},$$

where $\alpha(n)$ is the extremely slow growing inverse Ackermann function.

Definition 30. *The Ackermann function, first used by Wilhelm Ackermann [24], is defined as follows.*

$$A(m, n) = \begin{cases} n + 1 & \text{if } m = 0 \\ A(m - 1, 1) & \text{if } m > 0 \text{ and } n = 0 \\ A(m - 1, A(m, n - 1)) & \text{if } m > 0 \text{ and } n > 0. \end{cases}$$

The inverse Ackermann function, denoted by $\alpha(n)$, is the minimum r such that $A(r, r) \geq n$. \diamond

For our proof, we require the definition of Davenport-Schinzel sequences, which will help us bound the number of alternations between pairs of paths.

Definition 31. *Given a finite set of symbols \mathcal{X} , a sequence $U = (u_1, u_2, u_3, \dots, u_n)$ is a Davenport-Schinzel sequence of order s if it satisfies the following properties:*

- Each u_i (for $i \in [n]$) is a symbol coming from \mathcal{X} .
- No two consecutive symbols in the sequence U are the same.
- If $x_1 \in \mathcal{X}$ and $x_2 \in \mathcal{X}$ are two distinct symbols, then U does not contain a subsequence of the form $(\dots, x_1, \dots, x_2, \dots, x_1, \dots, x_2, \dots)$ consisting of $s + 2$ alternations between x_1 and x_2 .

Note that an alternation-free sequence is a Davenport-Schinzel sequence of order 1. \diamond

Proof of Theorem 2. We adapt to our setting an argument due to Gusfield [2, Page 100]. Fix an integer $d > 1$ and consider only those graphs whose edge weights are polynomials of degree at most d . Let $f(n, \ell)$ be the maximum length of a sequence of shortest paths⁶, when the paths are restricted to have at most ℓ edges. Fix a sequence σ of paths. Let p be a path in σ . We may fix a vertex v in p such that v is the middle vertex of the path p . That is, p has at most $\lceil \ell/2 \rceil$

⁶Note that this sequence might have alternations. That is, a shortest path can occur more than once in this sequence (however, two consecutive shortest paths must be distinct).

edges from s to v , and at most $\lceil \ell/2 \rceil$ edges from v to t . Then, the number of such paths in σ that pass through v is at most $2f(n, \lceil \ell/2 \rceil)$. Accounting for all v , we obtain that there are at most $2nf(n, \lceil \ell/2 \rceil)$ paths in the sequence σ . Now, σ might have alternations. Thus, if N is the number of *distinct* paths in σ , then $N \leq 2nf(n, \lceil \ell/2 \rceil)$. Since the costs of these paths are polynomials of degree at most d in λ , two paths can alternate at most $d + 1$ times (the curves of two distinct degree d polynomials cannot intersect each other in more than d points). That is, σ is a Davenport-Schinzel sequence of order d with an alphabet of size N . Bounds known for Davenport-Schinzel sequences (see Matoušek [25, Page 173]) imply that the maximum length of such a sequence of shortest paths is at most $N2^{\alpha(N)^d}$ (for all large N).

Since $N \leq n^n$ (a coarse upper bound on the total number of paths in an n -vertex graph), we have $\alpha(N) \ll \alpha(n) + 2$. Thus,

$$\begin{aligned} f(n, \ell) &\leq N \cdot 2^{\alpha(N)^d} \\ &\leq 2nf(n, \lceil \ell/2 \rceil) \cdot 2^{(\alpha(n)+2)^d} \\ &\leq (2n)^{\log \ell} \cdot f(n, 1) \cdot \left(2^{(\alpha(n)+2)^d}\right)^{\log \ell}. \end{aligned}$$

Our theorem follows from this after substituting $\ell = n$ and $f(n, 1) \leq 1$. \square

Remarks: (i) Note that even though the edge weights are polynomials of degree d in the parameter λ , the upper bound is not significantly higher than that for $d = 1$. (ii) Since our proof counts each time a path reappears as a separate occurrence, the same upper bound holds even if we count the number of different paths *with multiplicity*. (iii) Our proof works for any family of functions \mathcal{F} which satisfy the following conditions (the family of polynomials clearly satisfies these conditions).

- 1) \mathcal{F} is closed under addition.
- 2) For all $f_1, f_2 \in \mathcal{F}$, the sign of $f_1 - f_2$ can change only a small number of times.

VI. GRAPHS WITH EDGE WEIGHTS AS LINEAR FORMS

In this section, we consider graphs with edge weights of the form

$$w_e(\lambda_1, \lambda_2, \lambda_3) = a_e \lambda_1 + b_e \lambda_2 + c_e \lambda_3.$$

The main result of this section is the following.

Theorem 3 (Upper bound with three-parameter linear edge weights).

$$\varphi^{\text{lin}(3)}(n) = n^{(\log n)^2 + O(\log n)}.$$

Before we jump in to our proof, we note that if the edge weights have the form $a_e \lambda_1 + b_e \lambda_2$, then the upper bound is the same as the single parameter case. Consider the paths restricted to $\lambda_2 = -1$, and also the paths restricted to $\lambda_2 = 1$. Since all paths pass through the origin, it is easy to see that a path that forms a lower envelope without this restriction continues to be a part of the lower envelope for either $\lambda_2 = -1$ or for $\lambda_2 = 1$ (or for both). Therefore, we can match Gusfield's bound (up to a factor 2), which is for edge weights of the

form $a_e \lambda + b_e$. Thus, $\varphi^{\text{lin}(2)}(n) = n^{\log n + O(1)}$. Let us now commence our proof for the three parameter case.

Let G be a graph with edge weights of the form $a_e \lambda_1 + b_e \lambda_2 + c_e \lambda_3$. Note that the cost of an s - t path P in G has the form $\mathcal{C}(P)(\lambda_1, \lambda_2, \lambda_3) = a_P \lambda_1 + b_P \lambda_2 + c_P \lambda_3$. We may view these paths as vectors in \mathbb{R}^3 . That is, $\mathcal{C}(P)(\vec{\lambda}) = \vec{a}_P \cdot \vec{\lambda}$, where $\vec{\lambda} = (\lambda_1, \lambda_2, \lambda_3) \in \mathbb{R}^3$ and $\vec{a}_P = (a_P, b_P, c_P) \in \mathbb{R}^3$. We focus on the coefficient vectors \vec{a}_P . (The negative signs ensure that a longest coefficient vector corresponds to a shortest path.) Let $\text{conv}(G) \subseteq \mathbb{R}^3$ be the convex hull of the set $\{\vec{a}_P : P \text{ is an } s\text{-}t \text{ path in } G\}$. The crucial observation is that for each $\vec{\lambda} \in \mathbb{R}^3$, at least one of the vertices (vertex here means 0-dimensional face) of $\text{conv}(G)$ corresponds to a minimum cost path. Thus, the number of vertices of $\text{conv}(G)$ is an upper bound on $\varphi^{\text{lin}(3)}(G)$.

To upper bound the number of vertices of $\text{conv}(G)$, we proceed by induction roughly as in earlier proofs, but working in the language of polytopes. As before, fix ℓ and let $\text{conv}(G, \ell)$ be the convex hull of the set $\{\vec{a}_P : P \text{ is an } s\text{-}t \text{ path in } G \text{ with } \ell \text{ edges}\}$. We view every s - t path P with ℓ edges as consisting of an s - v path P_{sv} and a v - t path P_{vt} , each with about $\ell/2$ edges. Then the coefficient vector of the path P is the sum of the corresponding vectors of P_{sv} and P_{vt} , that is, $\vec{a}_P = \vec{a}_{P_{sv}} + \vec{a}_{P_{vt}}$. This naturally leads us to consider the two polytopes $\text{conv}(G_{sv}, \ell/2)$ and $\text{conv}(G_{vt}, \ell/2)$ generated by coefficients of the form $\vec{a}_{P_{sv}}$ and $\vec{a}_{P_{vt}}$, respectively. Thus, the convex hull of the coefficient vectors of paths bisected at v is the Minkowski sum $\text{conv}(G_{sv}, \ell/2) + \text{conv}(G_{vt}, \ell/2)$.

We show below how the number of vertices in such a Minkowski sum can be bounded. The final bound is obtained by running over all choices of v . Before presenting the formal proof, we show some facts about convex polytopes that we will need for our proof (we follow notation used by Michel Goemans in his lecture notes [26]).

A. Convex Polytopes and Minkowski Sums

Let $\mathcal{A} \subseteq \mathbb{R}^k$ be a non-empty convex polytope. The dimension of \mathcal{A} , denoted by $\dim(\mathcal{A})$, is the maximum number of affinely independent points in \mathcal{A} minus 1. A face of \mathcal{A} is a set of the form $\{\vec{x} \in \mathcal{A} : \vec{\lambda} \cdot \vec{x} = \beta\}$, where $\vec{\lambda} \cdot \vec{x} \geq \beta$ for all $\vec{x} \in \mathcal{A}$. Let $\mathcal{F}(\mathcal{A})$ be the set of faces of \mathcal{A} , and $\mathcal{F}^r(\mathcal{A})$ be the set of faces of \mathcal{A} of dimension r , where $0 \leq r \leq k$ (note that a face is also a polytope). Let $\varphi_r(\mathcal{A}) = |\mathcal{F}_r^r(\mathcal{A})|$ and $\varphi(\mathcal{A}) = |\mathcal{F}(\mathcal{A})|$.

The affine space spanned by the face F has the form $x_0 + \mathbb{S}_F$, where $x_0 \in F$ and \mathbb{S}_F is a subspace of dimension $\dim(F)$. We use Π_F to denote the orthogonal projection onto the space \mathbb{S}_F , which is the subspace of \mathbb{R}^d orthogonal to the subspace \mathbb{S}_F .

We say that a face F is non-trivial if $0 \neq \dim(F) \neq \dim(\mathcal{A})$. Then, every non-trivial face of \mathcal{A} is the set of solutions for a linear program of the form $\langle \min \vec{\lambda} \cdot \vec{x}; \text{subject to } \vec{x} \in \mathcal{A} \rangle$ with $\vec{\lambda} \neq 0$; we use $\mathfrak{L}(\mathcal{A}; \vec{\lambda})$ to refer to this set of solutions.

For polytopes \mathcal{A} and \mathcal{B} , let $\mathcal{A} + \mathcal{B}$ denote their Minkowski sum, that is, $\mathcal{A} + \mathcal{B} = \{\vec{x} + \vec{y} : \vec{x} \in \mathcal{A} \text{ and } \vec{y} \in \mathcal{B}\}$.

Proposition 32 (see [27, Proposition 2.1]). *Every face F of $\mathcal{A} + \mathcal{B}$ can be written uniquely as $F = F_A + F_B$, where $F_A \in \mathcal{F}(\mathcal{A})$ and $F_B \in \mathcal{F}(\mathcal{B})$. In particular, if*

$$F = \mathcal{L}(\mathcal{A} + \mathcal{B}; \vec{\lambda}), \quad (33)$$

then the unique decomposition is $F = \mathcal{L}(\mathcal{A}; \vec{\lambda}) + \mathcal{L}(\mathcal{B}; \vec{\lambda})$; in particular, $\mathcal{L}(\mathcal{A}; \vec{\lambda})$ and $\mathcal{L}(\mathcal{B}; \vec{\lambda})$ do not depend on the choice of $\vec{\lambda}$ in (33).

For a face F_A of \mathcal{A} , let $\mathcal{F}(\mathcal{A}, \mathcal{B}|F_A) = \{F \in \mathcal{F}(\mathcal{A} + \mathcal{B}) : F = F_A + F_B \text{ for some } F_B \in \mathcal{F}(\mathcal{B})\}$, and let $\varphi(\mathcal{A}, \mathcal{B}|F_A) = |\mathcal{F}(\mathcal{A}, \mathcal{B}|F_A)|$; that is, $\varphi(\mathcal{A}, \mathcal{B}|F_A)$ is the number of faces of $\mathcal{A} + \mathcal{B}$ whose unique decomposition (as in Proposition 32) involves F_A .

Lemma 34. *Suppose F_A is non-empty. Then, $\varphi(\mathcal{A}, \mathcal{B}|F_A) \leq \varphi(\Pi_{F_A}(\mathcal{B}))$.*

Proof. We will exhibit a one-to-one (injective) map from $\mathcal{F}(\mathcal{A}, \mathcal{B}|F_A)$ to $\mathcal{F}(\Pi_{F_A}(\mathcal{B}))$. To simplify notation for this proof, we will just write Π instead of Π_{F_A} . Fix a face $F = F_A + F_B$ of the polytope $\mathcal{A} + \mathcal{B}$. Let $\vec{\lambda}$ be such that $F = \mathcal{L}(\mathcal{A} + \mathcal{B}; \vec{\lambda})$, $F_A = \mathcal{L}(\mathcal{A}; \vec{\lambda})$ and $F_B = \mathcal{L}(\mathcal{B}; \vec{\lambda})$; note that $\vec{\lambda} \in \mathcal{S}_{F_A}^\perp$.

Claim 35. $\Pi(F_B) = \Pi(\mathcal{L}(\mathcal{B}; \vec{\lambda})) = \mathcal{L}(\Pi(\mathcal{B}); \vec{\lambda})$; in particular, $\Pi(F_A)$ is a face of $\Pi(\mathcal{B})$.

Proof. The first equality follows directly by our choice of $\vec{\lambda}$ above. Let $\beta = \vec{\lambda} \cdot \vec{x}$ for $\vec{x} \in F_B$. To justify the second equality, first observe that for all \vec{x} , we have

$$\vec{\lambda} \cdot \vec{x} = \vec{\lambda} \cdot \Pi(\vec{x}) \quad (\text{since } \vec{\lambda} \in \mathcal{S}_{F_A}^\perp). \quad (36)$$

Then for all $\vec{x} \in \mathcal{B}$, we have $\vec{\lambda} \cdot \Pi(\vec{x}) = \vec{\lambda} \cdot \vec{x} \geq \beta$. Also, if $\vec{x} \in \mathcal{L}(\mathcal{A}; \vec{\lambda})$, then $\vec{\lambda} \cdot \Pi(\vec{x}) = \beta$. Hence, $\Pi(\mathcal{L}(\mathcal{A}; \vec{\lambda})) \subseteq \mathcal{L}(\Pi(\mathcal{B}); \vec{\lambda})$. Next suppose $\vec{y} \in \mathcal{L}(\Pi(\mathcal{B}); \vec{\lambda})$, then $\vec{y} = \Pi(\vec{z})$ for some $\vec{z} \in \mathcal{B}$. By (36), we have $\vec{\lambda} \cdot \vec{z} = \vec{\lambda} \cdot \vec{y} = \beta$. Thus, $\vec{z} \in \mathcal{L}(\mathcal{B}; \vec{\lambda})$ and $\vec{y} \in \Pi(\mathcal{L}(\mathcal{B}; \vec{\lambda}))$. Hence, $\mathcal{L}(\Pi(\mathcal{B}); \vec{\lambda}) \subseteq \Pi(\mathcal{L}(\mathcal{A}; \vec{\lambda}))$. \square

Consider the map $\eta : \mathcal{F}(\mathcal{A}, \mathcal{B}|F_A) \rightarrow \mathcal{F}(\Pi(\mathcal{B}))$, which maps the face $F = F_A + F_B$ to the face $\Pi(F_B)$. By the above claim, this map is well-defined.

Claim 37. η is one-to-one.

Proof. Suppose $\eta(F_A + F_B) = \eta(F_A + F'_B)$. Then, by the definition of η , we have $\Pi(F_B) = \Pi(F'_B)$. Suppose $F_A + F_B = \mathcal{L}(\mathcal{A} + \mathcal{B}; \vec{\lambda})$ and $\vec{\lambda} \cdot \vec{x} = \beta$ for all $\vec{x} \in F_B$. In particular, $\forall \vec{z} \in \Pi(F_B)$, we have $\vec{\lambda} \cdot \vec{z} = \beta$. Then, for all $\vec{y} \in F'_B$, we have

$$\begin{aligned} \vec{\lambda} \cdot \vec{y} &= \vec{\lambda} \cdot \Pi(\vec{y}) && (\text{by (36)}) \\ &= \beta. && (\text{since } \Pi(\vec{y}) \in \Pi(F'_B) = \Pi(F_B)) \end{aligned}$$

Thus, $\vec{y} \in \mathcal{L}(\mathcal{B}; \vec{\lambda}) = F_B$. Hence, $F'_B \subseteq F_B$. Similarly, $F_B \subseteq F'_B$. \square

This completes the proof of Lemma 34. \square

B. An Upper Bound on the Number of Faces

We now formally carry out the inductive argument outlined earlier. The reason our proof works in three dimensions but not in higher dimensions is as follows. In the three-dimensional world, the number of faces of various dimensions (vertices, edges and faces) in a polytope are within small constant factors of each other, but there is no such bound in higher dimensions.

Proof of Theorem 3. Let \mathcal{P} be a convex polytope in three dimensions. Since each vertex in \mathcal{P} is connected to at least three other vertices by edges,

$$3\varphi_0(\mathcal{P}) \leq 2\varphi_1(\mathcal{P}). \quad (38)$$

Similarly, since each face in the polytope is adjacent to at least three other faces by edges,

$$3\varphi_2(\mathcal{P}) \leq 2\varphi_1(\mathcal{P}). \quad (39)$$

Combining Euler's formula, $\varphi_0(\mathcal{P}) - \varphi_1(\mathcal{P}) + \varphi_2(\mathcal{P}) = 2$, with (38) and (39), we obtain

$$2\varphi_0(\mathcal{P}) \leq \varphi_1(\mathcal{P}) + \varphi_2(\mathcal{P}) \leq 5\varphi_0(\mathcal{P}). \quad (40)$$

We now set up some notation in order to state a key recurrence. Let $r \in \{0, 1, 2, 3\}$ denote the dimension of a face. Let $\varphi_r^{\text{lin}(3)}(n, \ell) = \max_G \varphi_r(\mathbf{conv}(G, \ell))$ and $\varphi^{\text{lin}(3)}(n, \ell) = \max_G \varphi(\mathbf{conv}(G, \ell))$, where G varies over all graphs on n vertices. Let $\varphi_{1,2}^{\text{lin}(3)}(n, \ell) = \varphi_1^{\text{lin}(3)}(n, \ell) + \varphi_2^{\text{lin}(3)}(n, \ell)$. Similarly, for a convex polytope \mathcal{P} , let $\varphi_{1,2}(\mathcal{P}) = \varphi_1(\mathcal{P}) + \varphi_2(\mathcal{P})$.

We make the following claim, the proof of which can be found in the full version [23] of the paper.

Claim 41.

$$\varphi_{1,2}^{\text{lin}(3)}(n, \ell) \leq 5n \cdot \varphi^{\text{lin}(2)}(n) \cdot \varphi_{1,2}^{\text{lin}(3)}(n, \ell/2).$$

Using this claim, we are now ready to complete our proof. We have $\varphi^{\text{lin}(2)}(n) \leq n^{\log n + O(1)}$ (Gusfield's bound). Thus,

$$\begin{aligned} \frac{2}{3}\varphi^{\text{lin}(3)}(n, \ell) &\leq \varphi_{1,2}^{\text{lin}(3)}(n, \ell) \\ &\leq \left(5n \cdot n^{\log n + O(1)}\right) \cdot \varphi_{1,2}^{\text{lin}(3)}(n, \ell/2) \\ &\leq \left(5n \cdot n^{\log n + O(1)}\right)^{\log \ell} \\ &\leq n^{\log n \log \ell + O(\log \ell)} \\ \varphi^{\text{lin}(3)}(n, n) &\leq n^{(\log n)^2 + O(\log n)}. \end{aligned}$$

This completes the proof of Theorem 3. \square

VII. ALTERNATION-FREE SEQUENCES: LOWER BOUNDS

Theorem 6 (Alternation-free sequences of paths in layered graphs). *Let $G[n, m]$ be the layered graph with $m = 2^k$ layers, each layer consisting of n vertices (s is connected to all vertices in the first layer and t is connected to all vertices in the last layer). Let $f(n, m)$ be the length of a longest alternation-free sequence paths in $G[n, m]$; let $f^{\text{pl}}(n, m)$ be the length of a longest alternation-free sequence of paths in a planar subgraph of $G[n, m]$. Then, we have the following.*

$$(i) f(n, 2^k) \geq n^k, (ii) f^{\text{pl}}(n, (n-1)2^k) \geq n^k.$$

We construct non-planar and planar graphs with lengthy alternation-free sequences of paths. Our graphs are inspired by earlier examples of non-planar graphs with alternation-free sequences of paths ([2], [3]). For our construction, we introduce a concept of alternation-free sequences of *words* (in the case of planar graphs, these words will be binary strings), where each word corresponds to a path. It so turns out that these words, when arranged in a standard lexicographic order, correspond to an alternation-free sequence of paths. In Subsection VII-A, we will formally build this connection between paths and words. In Subsection VII-B and Subsection VII-C, we will use this connection to show Theorem 6.

Corollary 42 (of Theorem 6).

$$(i) f(n, n) \geq n^{\log n}, (ii) f^{\text{pl}}(n, n^2 - n) \geq n^{\log n}.$$

Viewing the construction behind the proof of Theorem 6 this from the treewidth lens, we have the following connection between pathwidth and alternation-free sequences in certain graphs.

Corollary 43 (of Theorem 6). *For every $n, \kappa \geq 4$, there is a graph on $n\kappa$ vertices of pathwidth κ for which there exists an alternation-free sequence of length $n^{\Omega(\log \kappa)}$.*

A. Alternation-free Sequences of Words

We first present an alternation-free sequence of paths for non-planar graphs, and then refine it to obtain another for a related planar graph. Consider the graph $G[n, m]$ with vertex set

$$V = \{(i, j) : i = 0, 1, \dots, n-1; j = 0, 1, \dots, m-1\} \cup \{s, t\}.$$

We partition $V \setminus \{s, t\}$ into layers L_0, L_1, \dots, L_{m-1} of n vertices each, where the j -th layer is

$$L_j = \{(i, j) : i = 0, 1, \dots, n-1\}.$$

There are edges from vertex s to all vertices in L_0 , and from all vertices in L_{m-1} to t . The remaining edges connect vertices in one layer to the vertices in the next. We will have two version of the graph: a non-planar version and a planar version. Let $\mathbb{Z}_n = \{0, 1, \dots, n-1\}$, with addition performed modulo n . In the non-planar version, we add all edges from a layer to the next. We refer to the resulting graph as $G^{\text{nppl}}[n, m]$:

$$E(G^{\text{nppl}}) = (\{s\} \times L_0) \cup (L_0 \times L_1) \cup \dots \cup (L_{m-1} \times \{t\}).$$

Thus, for $j = 0, 1, \dots, m-2$, the subgraph induced by $L_j \cup L_{j+1}$ is a complete bipartite graph. In the planar version, we connect a vertex in layer j to two vertices in layer $j+1$. We refer to the resulting graph as $G^{\text{pl}}[n, m]$:

$$E(G^{\text{pl}}) = \{((i, j), (i+b \bmod n, j+1)) : b \in \{0, 1\}, \\ i = 0, 1, \dots, n-1, \text{ and } j = 0, 1, \dots, m-2\}.$$

One may imagine that G^{pl} is drawn on the surface of a cylinder instead of the surface of a plane. This means that the $(n-1)$ -th vertex in layer L_j goes around the surface of the cylinder to the 0-th vertex in layer L_{j+1} (Figure 6). This does not affect

the planarity of the graph (a graph that can be drawn on the surface of a cylinder can also be drawn on a planar surface). In G^{nppl} , we may encode s - t paths by words in \mathbb{Z}_n^m : the word $\sigma = (\sigma_0, \sigma_1, \dots, \sigma_{m-1}) \in \mathbb{Z}_n^m$ corresponds to the path

$$p_\sigma = (s, (i_0, 0), (i_1, 1), \dots, (i_{m-1}, m-1), t),$$

where $i_0 = \sigma_0$, and $i_{j+1} = i_j + \sigma_{j+1} \bmod n$, for $j = 0, 1, \dots, m-2$. Similarly, we associate words $\tau \in \{0, 1\}^m$ with paths p_τ in G^{pl} . We define alternation-free sequences of words, and observe that the corresponding paths are alternation-free. By showing long alternation-free sequences of words, we establish the existence of long alternation-free sequences of paths.

Definition 44 (Word). *Let \mathbb{Z}_n denote the set $\{0, 1, 2, \dots, n-1\}$ where addition is performed modulo n . Let \mathbb{Z}_n^m denote the set of words over \mathbb{Z}_n of length m . For a word $\sigma \in \mathbb{Z}_n^m$ and $i \in \{0, 1, \dots, m-1\}$, let $\sigma[i]$ denote the i -th element of σ ; let $\sigma[i : j]$ denote the subword $(\sigma[i], \sigma[i+1], \dots, \sigma[j-1])$. For $\sigma \in \mathbb{Z}_n^k$, let $|\sigma|_1$ denote the sum (in \mathbb{Z}_n) of its elements. That is, $|\sigma|_1 = \sum_{i=0}^{k-1} \sigma[i] \bmod n$. Given a word $\sigma \in \mathbb{Z}_n^m$ and $j \in \mathbb{Z}_n$, let $\sigma \downarrow j$ be the word $\mu \in \mathbb{Z}_n^{2m}$ obtained from σ by inserting j after each symbol of σ . That is, if $\mu = \sigma \downarrow j$, then $\mu[2i] = \sigma[i]$ and $\mu[2i+1] = j$, for $i = 0, 1, \dots, m-1$. Let $S \in (\mathbb{Z}_n^m)^\ell$ be a sequence of words (each word is in \mathbb{Z}_n^m), and $S \downarrow j = (\sigma \downarrow j : \sigma \in S)$ be the sequence obtained after performing such an insertion on every word of S . For instance, if $\sigma = (7 \ 2 \ 6 \ 2)$, then $\sigma \downarrow 3 = (7 \ 3 \ 2 \ 3 \ 6 \ 3 \ 2 \ 3)$. \diamond*

Definition 45 (Alternation-free sequence of words). *Let S be sequence of ℓ words, each from \mathbb{Z}_n^m . That is, $S = (\sigma_1, \sigma_2, \dots, \sigma_{\ell-1}) \in (\mathbb{Z}_n^m)^\ell$. We say that S has an alternation at (a, b, c) between (u, v) , where $0 \leq a < b < c \leq \ell-1$ and $0 \leq u < v \leq m$, if*

- $|\sigma_a[0 : u]|_1 = |\sigma_b[0 : u]|_1 = |\sigma_c[0 : u]|_1$;
- $v = m$ or $(|\sigma_a[0 : v]|_1 = |\sigma_b[0 : v]|_1 = |\sigma_c[0 : v]|_1)$;
- $\sigma_a[u : v] = \sigma_c[u : v] \neq \sigma_b[u : v]$.

Note that in every such alternation, either $v = m$ or $v - u \geq 2$. If S has no alternation, then we say that S is an alternation-free sequence of words. \diamond

Proposition 46 (Paths from sequences). *If $S = (\sigma_i : i = 0, 1, \dots, \ell-1) \in (\mathbb{Z}_n^m)^\ell$ is an alternation-free sequence of words, then $(p_{\sigma_i} : i = 0, 1, \dots, \ell-1)$ is an alternation-free sequence of paths in G^{nppl} . Similarly, if $T = (\tau_i : i = 0, 1, \dots, \ell-1) \in (\{0, 1\}^m)^\ell$ is an alternation-free sequence of words, then $(p_{\tau_i} : i = 0, 1, \dots, \ell-1)$ is an alternation-free sequence of paths in G^{pl} .*

Proof. Straightforward. Note that the case $v = m$ in the second condition of Definition 45 is used to verify that there is no alternation involving pairs of vertices of the form (u, t) . \square

Thus, we can now focus on creating alternation-free sequences of words.

B. Constructing Alternation-free Sequences of Words

In this section, we will construct two alternation-free sequences X and \hat{X} (each of length n^ℓ) over \mathbb{Z}_n and $\{0, 1\}$




Figure 6: The graph G^{p1} for $n = 6, m = 8$. s is connected to all vertices in the top layer and t is connected to all vertices in the bottom layer. The path corresponds to the word $(0) \downarrow 2 \downarrow 4 \downarrow 1$.

respectively. We first describe X . The i -th word ($i = 0, 1, \dots, n^\ell - 1$) of X is given by $X[i] = (b_0) \downarrow b_1 \downarrow \dots \downarrow b_{\ell-1}$, where $(i)_n = \sum_{j=0}^{\ell-1} b_j n^j$ is the base n representation of i . For example, suppose $n = 6, m = 8$ and $i = 372$. Then $X[372] = (0) \downarrow 2 \downarrow 4 \downarrow 1 = (0 \ 1 \ 4 \ 1 \ 2 \ 1 \ 4 \ 1)$ because the decimal number 372 is equal to 1420 in base 6.

Binary alternation-free sequences can be viewed as a composition of words over \mathbb{Z}_n , where we map $i \in \mathbb{Z}_n$ to the binary word $\hat{i} = 1^i 0^{n-1-i} \in \{0, 1\}^{n-1}$. Thus $\hat{X}[i]$ is constructed exactly like $X[i]$, but it is represented differently (as a binary word of length $(n-1)2^{\ell-1}$ bits). Continuing with the example in the previous paragraph, we have $\hat{X}[372] = (00000 \ 10000 \ 11110 \ 10000 \ 10000 \ 10000 \ 11110 \ 10000)$. See Figure 6 for a 3D depiction of this path on the surface of a cylinder. Now we will show that X and \hat{X} are alternation-free.

Lemma 47. Suppose $S \in (\mathbb{Z}_n^m)^\ell$ is an alternation-free sequence of ℓ words in \mathbb{Z}_n^m . Then,

- (a) for all $j \in \mathbb{Z}_n$, $S \downarrow j$ is an alternation-free sequence of ℓ words in \mathbb{Z}_n^{2m} ;
- (b) $T = (S \downarrow 0) \circ (S \downarrow 1) \circ \dots \circ (S \downarrow (n-1))$ is an alternation-free sequence of $n\ell$ words, where each word is in \mathbb{Z}_n^{2m} .

Proof. For part (a), note that if $S \downarrow j$ has an alternation at (a, b, c) ($0 \leq a < b < c \leq \ell - 1$) between (u, v) ($0 \leq u < v \leq 2m$), then it is easy to see that S itself has an alternation at (a, b, c) , between $(\lceil u/2 \rceil, \lceil v/2 \rceil)$. Since S is alternation-free, so is $S \downarrow j$.

For part (b), we use part (a). Suppose T has an alternation at (a, b, c) ($0 \leq a < b < c \leq n\ell - 1$) between (u, v) ($0 \leq u < v \leq 2m$). If σ_a and σ_c have the same symbol in their odd positions⁸ then σ_a, σ_b and σ_c all come from a common segment of T of the form $S \downarrow j$. By part (a), the sequence $S \downarrow j$ is alternation-free. So T has no alternation at (a, b, c) between (u, v) .

On the other hand, suppose σ_a and σ_c have different symbols in their odd positions. Since $\sigma_a[v-1] = \sigma_c[v-1]$, we conclude that v is odd. In particular, $v \neq 2m$ and thus $v-u \geq 2$ (as observed at the end of Definition 45). This means that the interval $\{u, u+1, \dots, v-1\}$ includes an odd number. Hence $\sigma_a[u : v] \neq \sigma_c[u : v]$, and there is no alternation at (a, b, c) between (u, v) . This completes the proof. \square

Theorem 48. For all $\ell \geq 1$, there is an alternation-free sequence T of n^ℓ words in $\mathbb{Z}_n^{2^\ell}$.

Proof. We will use Lemma 47 and induction on ℓ . For $\ell = 1$, the alternation-free sequence is simply $T = (0, 1, 2, \dots, n-1)$, which we think of as a sequence of n words, where each word has one symbol. Suppose $\ell > 1$. Let S be sequence of n^ℓ words in $\mathbb{Z}_n^{2^{\ell-1}}$. Consider the sequence

$$T = (S \downarrow 0) \circ (S \downarrow 1) \circ \dots \circ (S \downarrow (n-1)).$$

⁷Here \circ represents concatenation of sequences.

⁸Recall that σ has j in all its odd positions, for every $\sigma \in S \downarrow j$.

By Lemma 47, T is an alternation-free sequence of $n \cdot n^\ell = n^{\ell+1}$ words in $\mathbb{Z}_n^{2 \cdot 2^{\ell-1}} = \mathbb{Z}_n^{2^\ell}$. \square

Note that T is equal to the X that we had described earlier.

C. Constructing Alternation-free Sequences of Binary Words

Consider the following *unary* encoding, where we map $i \in \mathbb{Z}_n$ to the binary word $\hat{i} = 1^i 0^{n-1-i} \in \{0, 1\}^{n-1}$. Let $\hat{\mathbb{Z}}_n = \{\hat{0}, \hat{1}, \dots, \widehat{n-1}\}$. Thus, words in $\hat{\mathbb{Z}}_n^m$ over this alphabet consist of m symbols, each of which is a binary word of $n-1$ bits. We view such a word as a binary string of length $m(n-1)$ by concatenating the m symbols. Now we will show that the resulting sequence of binary strings is alternation-free.

Lemma 49. *Suppose $\hat{S} \in (\hat{\mathbb{Z}}_n)^\ell$ is alternation-free sequence of ℓ binary words of length $m(n-1)$ each. Then,*

- (a) *for all $\hat{j} \in \hat{\mathbb{Z}}_n$, $\hat{S} \downarrow \hat{j}$ is an alternation-free sequence of ℓ words in $\{0, 1\}^{2m(n-1)}$;*
- (b) *$\hat{T} = (\hat{S} \downarrow \hat{0}) \circ (\hat{S} \downarrow \hat{1}) \circ \dots \circ (\hat{S} \downarrow \widehat{n-1})$ is an alternation-free sequence of $n\ell$ words in $\{0, 1\}^{2m(n-1)}$.*

Proof. For part (a), consider $\hat{S} \downarrow \hat{j}$. A word in this sequence consists of blocks of $n-1$ symbols, where each block can be thought of as an element of $\hat{\mathbb{Z}}_n$. In particular, all the odd numbered blocks of $\hat{S} \downarrow \hat{j}$ comprise of the word \hat{j} . Since the symbols from these odd blocks make the same contribution to the prefix sums of all words, we can suppress them and conclude that $\hat{S} \downarrow \hat{j}$ is alternation-free because \hat{S} is known to be alternation-free. We now make this idea more precise.

Suppose $\hat{S} \downarrow \hat{j} = (\sigma_i : i = 0, 1, \dots, \ell-1)$ has an alternation at (a, b, c) ($0 \leq a < b < c \leq \ell-1$) between (u, v) ($0 \leq u < v \leq 2m(n-1)$). First, let us handle the boundary case $v = 2m(n-1)$. Then u cannot be a location in the last block, for it is an odd numbered block and the entire block is identical in all words in $\hat{S} \downarrow \hat{j}$. Let $u = q(n-1) + r$, where $r = u \bmod n-1$ and $q < 2m-1$. We conclude that \hat{S} has an alternation at (a, b, c) between $(\lceil q/2 \rceil (n-1) + r, m(n-1))$, contradicting the fact that \hat{S} is alternation-free.

So we may assume that $v < 2m(n-1)$ and $v-u \geq 2$. We may also assume that (u, v) has been chosen so that $v-u$ is minimal. This implies that $\sigma_a[u] = \sigma_c[u] \neq \sigma_b[u]$, and similarly that $\sigma_a[v-1] = \sigma_c[v-1] \neq \sigma_b[v-1]$. In particular, both u and $v-1$ are indices in even numbered blocks. Let $u = q(n-1) + r$, $v-1 = q'(n-1) + r'$, where $r = u \bmod n-1$ and $r' = v-1 \bmod n-1$. Then, q and q' are even. We conclude that \hat{S} has an alternation at (a, b, c) between $((q/2)(n-1) + r, (q'/2)(n-1) + r')$, contradicting the fact that \hat{S} is alternation-free. This establishes part (a).

For part (b), suppose \hat{T} has an alternation at (a, b, c) ($0 \leq a < b < c \leq n\ell-1$) between (u, v) ($0 \leq u < v \leq 2m(n-1)$); assume that (u, v) has been chosen so that $v-u$ is minimal. Let $\sigma_a \in \hat{S} \downarrow \hat{j}_a$, $\sigma_c \in \hat{S} \downarrow \hat{j}_c$ for some $\hat{j}_a, \hat{j}_c \in \hat{\mathbb{Z}}_n$. If $\hat{j}_a = \hat{j}_c$, then part (a) gives us the necessary contradiction. So we may assume that $\hat{j}_a \neq \hat{j}_c$. Note that the contents of the odd numbered blocks are *monotonically non-decreasing* in $\hat{\mathbb{Z}}_n$. Thus if an entire odd numbered block lies in the range $\{u, u+1, \dots, v-1\}$, then $\sigma_a[u : v] \neq \sigma_c[u : v]$, and there is

no alternation. So we may assume that no entire odd numbered block lies in the range $\{u, u+1, \dots, v-1\}$.

To complete the proof, we now use two crucial facts about our encoding. First, in the sequence of words $(\hat{0}, \hat{1}, \dots, \widehat{n-1})$, the bit at position i (for $i = 0, 1, \dots, n-1$) starts at 0 and flips to 1, never to flip back to 0 again. This implies that both u and $v-1$ lie in even numbered blocks. Since no entire odd numbered block lies in the range $\{u, u+1, \dots, v-1\}$, this means that u and $v-1$ lie within the same even numbered block. Second, our encoding has the property that if for two words $\hat{j}_1, \hat{j}_2 \in \hat{\mathbb{Z}}_n$, we have $|\hat{j}_1[u : v]|_1 = |\hat{j}_2[u : v]|_1$, then $\hat{j}_1[u : v] = \hat{j}_2[u : v]$. This implies that $\sigma_a[u : v] = \sigma_b[u : v] = \sigma_c[u : v]$; so there is no alternation. This establishes part (b) and completes the proof. \square

Note that \hat{T} is equal to the \hat{X} that we had described earlier.

VIII. THE PRAM LOWER BOUNDS

Mulmuley & Shah's [3] Theorem 1.4 states the following.

Claim 50. *The Shortest Path Problem cannot be computed in $o(\log n)$ steps on an unbounded fan-in PRAM without bit operations using $\text{poly}(n)$ processors, even if the weights on the edges are restricted to have bit-lengths $O((\log n)^2)$.*

A more precise statement of their result (see also Theorem 4.2.1 of Pradyut Shah's PhD thesis [28]) is the following: There exist constants $\alpha > 0$ and $\epsilon > 0$, and an explicitly described family of weighted graphs G_n (G_n has n vertices and weights that are $O((\log n)^2)$ bits long), such that for infinitely many n , every algorithm on an unbounded fan-in PRAM without bit operations with at most n^α processors requires at least $\epsilon \log n$ steps to compute the shortest s - t path in G_n . (Their proof yields a constant $\alpha < 1$.)

Our proof of Claim 4, like Mulmuley & Shah's proof of the corresponding theorem [3, Theorem 1.4], is based on the following (see [3, Theorem 1.1]).

Theorem 51. *Let $\Phi(n, \beta(n))$ be the parametric complexity of any homogeneous optimization problem where n denotes the input cardinality and $\beta(n)$ the bit-size of the parameters. Then the decision version of the problem cannot be solved in the PRAM model without bit operations in $o(\sqrt{\log \Phi(n, \beta(n))})$ time using $2^{\sqrt{\log \Phi(n, \beta(n))}}$ processors even if we restrict every numeric parameter in the input to size $O(\beta(n))$.*

A version of Theorem 51 for *bounded fan-in* PRAMs is established in Mulmuley [29, Theorem 3.3]; Mulmuley & Shah [3] state that this theorem is also applicable to unbounded fan-in PRAMs. Unfortunately, no formal justification of this latter claim seems to be available in the literature (see Shah [28, Page 36] for an informal justification).

IX. THE WEIGHTED GRAPH MATCHING PROBLEM

In this section, we will see that the shortest path problem reduces to the WEIGHTED GRAPH MATCHING problem. More precisely, given a directed acyclic graph G with non-negative edge weights and two special vertices s and t , we will show

that a minimum weight perfect matching in G' (a slight modification G) can be used to compute an s - t shortest path in G . We now describe how to construct G' from G .

For every vertex $v \in V(G) \setminus \{s, t\}$, replace v by two vertices v_{in} and v_{out} and add a 0-weight edge $(v_{\text{in}}, v_{\text{out}})$ between them. This edge is the only outgoing edge from v_{in} and the only incoming edge to v_{out} . The in-neighbours of v_{in} are the in-neighbours of v , and the out-neighbours of v_{out} are the out-neighbours of v . Call this new graph G' .

It is easy to see that G' always has a perfect matching. Let M be a minimum weight perfect matching in G' . Let v_{in}^1 be the partner of s in M . This means that the edge $(v_{\text{in}}^1, v_{\text{out}}^1)$ is not in M . Now, let v_{in}^2 be the partner of v_{out}^1 in M . This means that the edge $(v_{\text{in}}^2, v_{\text{out}}^2)$ is not in M . Carrying this argument forward, we obtain that the edge $(v_{\text{in}}^r, v_{\text{out}}^r)$ is not in M (for some r), where v_{out}^r is the partner of t in M .

Note that the path $(s, v_{\text{in}}^1, v_{\text{out}}^1, \dots, v_{\text{in}}^r, v_{\text{out}}^r, t)$ is an s - t shortest path in G (otherwise, a path of lower weight can be used to obtain a matching of lower weight than M). In fact, the cost of this path is precisely the weight of M , as all the other edges of M are of weight 0.

X. CONCLUSION

Our ultimate goal is to understand how parametric shortest path complexity changes with the topology of the graph. One of the reasons Nikolova's conjecture was open is that planar graphs have a small (linear) number of edges, which leads one to (falsely) believe that there are not enough edge weights to assign for the graph to have high parametric complexity. Thus, the number of edges is not the right measure to characterize parametric complexity.

Since graphs with large treewidth have superpolynomial parametric complexity and graphs with constant treewidth have polynomial parametric complexity, treewidth seems to be the right measure. However, an unexplored gap still remains. The following conjecture, in particular, is interesting.

Conjecture 52. *For every sufficiently large n, k , there is an n -vertex graph of treewidth k with parametric complexity $n^{\Omega(\log k)}$.*

This seems plausible for two reasons: (i) there is a graph on n vertices of pathwidth k having parametric complexity $n^{\Omega(\log k - \log \log n)}$; (ii) there is an a graph on nk vertices of pathwidth k for which there exists an alternation-free sequence of length $n^{\Omega(\log k)}$ (Corollary 43). These results suggest that we might be very close to resolving Conjecture 52.

We used alternation-free sequences as a combinatorial way to view parametric shortest paths. Although Kuchlbauer's counterexample [19, Example 3.11] shows that there are infeasible alternation-free sequences, the following question is interesting: what is the worst-case minimum number of paths that we need to delete from an infeasible alternation-free sequence of paths so that the sequence becomes feasible?

Conjecture 53. *There exists a universal constant γ ($0 < \gamma < 1$) such that, for every graph G , if the length of the*

longest alternation-free sequence of paths in G is L , then the parametric complexity of G is at least L^γ .

If this is true, then alternation-free sequences bound parametric shortest paths from both below and above. More generally, the structure of a graph would completely determine its parametric complexity. However, all known methods for constructing graphs with large parametric complexity do not use the fact that the shortest paths form an alternation-free sequence, and thus it might require considerable insight to resolve Conjecture 53.

The total number of different paths in a directed acyclic graph on n vertices can be as high as $\exp(n)$. The original problem on parametric shortest paths considers linear edge weights in one variable, which yields shortest path complexity $n^{\Theta(\log n)}$. We generalized the edge weights to polynomials in one variable (Theorem 2), and to linear forms in three variables (Theorem 3). In both cases, the upper bounds are only slightly higher than those for univariate linear edge weights, and nowhere near $\exp(n)$. It is thus natural to consider a further generalization to multivariate polynomial edge weights.

Conjecture 54. *Let G be an n -vertex graph whose edge weights are polynomials in k variables of degree at most d . Then the parametric complexity of G is at most $n^{(\log n)^\epsilon}$, where $\epsilon = \text{poly}(k, d)$.*

Finally, it is interesting to explore other function families, and to see if there is a sequence of graphs $\{G_n\}$, where G_n has n vertices, with a well-defined set of edge weight functions on \mathbb{R} such that the parametric complexity of G_n is $\exp(n)$.

Acknowledgments

We thank Jannik Matuschke for introducing us to the problem and for several subsequent discussions, Tulasimohan Molli for helping us with the initial analysis of alternation-free sequences in planar graphs, Martina Kuchlbauer for sharing with us her example of infeasible alternation-free sequences, Suhail Sherif for helping us simulate the problem using a computer program, and Mark Wibrow for his 3D depiction of a cylinder in \LaTeX , which we used as a template for Figure 6. We also thank Vaishali Surianarayanan for suggesting that we consider multivariate linear forms, and Hariharan Narayanan for helpful discussions on convex polytopes. Finally, we thank the anonymous reviewers of prior versions of this paper for patiently going over the proofs, and helping us correct the errors and improve the overall presentation of the paper.

REFERENCES

- [1] P. J. Carstensen, "Complexity of some parametric integer and network programming problems," *Math. Program.*, vol. 26, no. 1, pp. 64–75, 1983. [Online]. Available: <https://doi.org/10.1007/BF02591893>
- [2] —, "The complexity of some problems in parametric linear and combinatorial programming," Ph.D. dissertation, University of Michigan, 1983.
- [3] K. Mulmuley and P. Shah, "A lower bound for the shortest path problem," *J. Comput. Syst. Sci.*, vol. 63, no. 2, pp. 253–267, 2001. [Online]. Available: <https://doi.org/10.1006/jcss.2001.1766>
- [4] E. Nikolova, "Strategic algorithms," Ph.D. dissertation, Massachusetts Institute of Technology, 2009.

- [5] D. M. Gusfield, "Sensitivity analysis for combinatorial optimization," Ph.D. dissertation, University of California, Berkeley, 1980.
- [6] —, "Parametric combinatorial computing and a problem of program module distribution," *Journal of the ACM (JACM)*, vol. 30, no. 3, pp. 551–563, 1983.
- [7] P. Kasteleyn, "Graph theory and crystal physics," *Graph theory and theoretical physics*, pp. 43–110, 1967.
- [8] L. Csanky, "Fast parallel matrix inversion algorithms," in *Foundations of Computer Science, 1975., 16th Annual Symposium on.* IEEE, 1975, pp. 11–12.
- [9] N. Anari and V. V. Vazirani, "Planar graph perfect matching is in NC," in *Proceedings of the 59th IEEE Annual Symposium on Foundations of Computer Science.* IEEE Computer Society, October 2018.
- [10] P. Sankowski, "NC algorithms for weighted planar perfect matching and related problems," in *45th International Colloquium on Automata, Languages, and Programming (ICALP 2018)*, ser. Leibniz International Proceedings in Informatics (LIPIcs), vol. 107. Dagstuhl, Germany: Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, 2018, pp. 97:1–97:16. [Online]. Available: <http://drops.dagstuhl.de/opus/volltexte/2018/9101>
- [11] C. Chekuri and J. Chuzhoy, "Polynomial bounds for the grid-minor theorem," *Journal of the ACM (JACM)*, vol. 63, no. 5, p. 40, 2016.
- [12] K. Gajjar, P. Misra, and J. Radhakrishnan, "Parametric Shortest Paths in Graphs of Bounded Treewidth," *Under preparation*, 2019.
- [13] E. Nikolova, J. A. Kelner, M. Brand, and M. Mitzenmacher, "Stochastic shortest paths via quasi-convex maximization," in *European Symposium on Algorithms.* Springer, 2006, pp. 552–563.
- [14] J. R. Correa, T. Harks, V. J. C. Kreuzen, and J. Matuschke, "Fare evasion in transit networks," *Operations Research*, vol. 65, no. 1, pp. 165–183, 2017. [Online]. Available: <https://doi.org/10.1287/opre.2016.1560>
- [15] S. Chakraborty, E. Fischer, O. Lachish, and R. Yuster, "Two-phase algorithms for the parametric shortest path problem," in *27th International Symposium on Theoretical Aspects of Computer Science-STACS 2010*, 2010, pp. 167–178.
- [16] J. Erickson, "Maximum flows and parametric shortest paths in planar graphs," in *Proceedings of the twenty-first annual ACM-SIAM symposium on Discrete Algorithms.* Society for Industrial and Applied Mathematics, 2010, pp. 794–804.
- [17] G. Borradaile and P. Klein, "An $O(n \log n)$ algorithm for maximum st-flow in a directed planar graph," *J. ACM*, vol. 56, no. 2, pp. 9:1–9:30, Apr. 2009. [Online]. Available: <http://doi.acm.org/10.1145/1502793.1502798>
- [18] R. M. Karp and J. B. Orlin, "Parametric shortest path algorithms with an application to cyclic staffing," *Discrete Applied Mathematics*, vol. 3, no. 1, pp. 37–45, 1981.
- [19] M. Kuchlbauer, "Parametric combinatorial optimization problems and their complexity," Master's thesis, Technische Universitt Mnchen, Munich, Bavaria, 80333, Germany, 2018.
- [20] G. C. Shephard, "Twenty problems on convex polyhedra part ii," *The Mathematical Gazette*, vol. 52, no. 382, pp. 359–367, 1968.
- [21] B. Chazelle, H. Edelsbrunner, and L. J. Guibas, "The complexity of cutting complexes," *Discrete & Computational Geometry*, vol. 4, no. 2, pp. 139–181, Mar 1989. [Online]. Available: <https://doi.org/10.1007/BF02187720>
- [22] J. C. Lagarias, Y. Luo, and A. Padrol, "Moser's shadow problem," *arXiv preprint arXiv:1310.4345*, 2013.
- [23] K. Gajjar and J. Radhakrishnan, "Parametric shortest paths in planar graphs," *Electronic Colloquium on Computational Complexity (ECCC)*, vol. 25, p. 211, 2018. [Online]. Available: <https://ecc.weizmann.ac.il/report/2018/211>
- [24] W. Ackermann, "Zum hilbertschen aufbau der reellen zahlen," *Mathematische Annalen*, vol. 99, no. 1, pp. 118–133, Dec 1928. [Online]. Available: <https://doi.org/10.1007/BF01459088>
- [25] J. Matoušek, *Lectures on discrete geometry.* Springer New York, 2002, vol. 212.
- [26] M. X. Goemans, "3. linear programming and polyhedral combinatorics," *Massachusetts Inst. Technol., Cambridge, MA, USA*, 2009.
- [27] K. Fukuda, "From the zonotope construction to the minkowski addition of convex polytopes," *Journal of Symbolic Computation*, vol. 38, no. 4, pp. 1261–1272, 2004.
- [28] P. Shah, "Lower bounds for parallel algorithms," Ph.D. dissertation, The University of Chicago, 2001.
- [29] K. Mulmuley, "Lower bounds in a parallel model without bit operations," *SIAM Journal on Computing*, vol. 28, no. 4, pp. 1460–1509, 1999.