Electronic Colloquium on Computational Complexity, Revision 1 of Report No. 211 (2018)

Parametric Shortest Paths in Planar Graphs

Kshitij Gajjar* and Jaikumar Radhakrishnan†

Tata Institute of Fundamental Research, Mumbai

February 4, 2019

Abstract

We construct a family of planar graphs ($G_n : n \geq 4$), where G_n has n vertices including a source vertex s and 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 $n^{\Omega(\log n)}$ break points. This shows that lower bounds obtained earlier by Carstensen (1983) and Mulmuley & Shah (2000) for general graphs also hold for planar graphs. A conjecture of Nikolova (2009) states that the number of break points in n-vertex planar graphs is bounded by a polynomial in n; our result refutes this conjecture.

Gusfield (1980) and Dean (2009) showed that the number of break points for every n-vertex graph is $n^{\log n + O(1)}$ assuming linear edge weights; we show that if the edge weights are allowed to vary as a polynomial of degree at most d, then the number of break points is $n^{\log n + (\alpha(n)+O(1))^d}$, where $\alpha(n)$ is the slowly growing inverse Ackermann function. This upper bound is obtained by appealing to bounds on the length of Davenport-Schinzel sequences.

1 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 $w_e(\lambda) = m_e\lambda + c_e$, for some real numbers m_e and c_e. The cost (also referred to as length) 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 $w_p(\lambda) = m_p\lambda + c_p$. The cost of the shortest s-t path is then given by

$$\mathcal{C}(\lambda) = \min_p w_p(\lambda),$$

where p ranges over all s-t paths; this function is the piece-wise linear *lower envelope* of the linear costs provided by the s-t paths. The main object of our investigation is the number of break

*kshitij.gajjar@tifr.res.in
†jaikumar@tifr.res.in

ISSN 1433-8092

points in this envelope, which 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 break points 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 $\varphi^{\mathrm{pl}}(n, \beta(n))$ be the parametric complexity when the graphs are restricted to be planar. We show the following.

Theorem 1 (Main result)**.** $\varphi^{\mathrm{pl}}(n, O((\log n)^3)) = n^{\Omega(\log n)}$.

Before this work similar results were known for general graphs. Carstensen [Car83a, Car83b] showed that $\varphi(n, \infty) = n^{\Omega(\log n)}$; her result was simplified and extended by Mulmuley & Shah [MS01], who showed that $\varphi(n, O((\log n)^3)) = n^{\Omega(\log n)}$. Carstensen [Car83b, Page 100] also presented a matching upper bound argument, $\varphi(n, \infty) = n^{\log n + O(1)}$, which she attributed to Daniel Gusfield [Gus80, Gus83] (a similar argument, attributed to Brian Dean, was presented by Nikolova [Nik09, Page 86]). For *planar graphs,* however, the complexity remained open [Nik09, Conjecture 6.1.6].

Conjecture 2 (Nikolova [Nik09])**.** $\varphi^{\mathrm{pl}}(n, \infty) = n^{O(1)}$.

Our main result provides a strong (with bit length $O((\log n)^3)$) refutation of this conjecture.

1.1 Parametric complexity and treewidth

We display a close relationship between the treewidth of a graph and its parametric complexity. We show this through two implications (Corollary 3, Corollary 4) of our main result (Theorem 1). The first implication deals with grid graphs. Since every planar graph can be embedded in a grid graph with a small blow-up in size, our lower bound holds for $n \times n$ grid graphs as well. Let $\varphi^{\mathrm{gr}}(n, n, \beta(n))$ be the parametric complexity when the graphs are restricted to be $n \times n$ grid graphs.

Corollary 3 (Square grid graphs)**.** $\varphi^{\mathrm{gr}}(n, n, O((\log n)^3)) = n^{\Omega(\log n)}$.

We also explore the parametric complexity of $k \times n$ grid graphs when $k \ll n$ (see Appendix B). The second implication deals with graphs of large treewidth. Using Chekuri & Chuzhoy's result [CC16] that every graph of treewidth k (for large enough k) has an $\Omega(k^{1/99}) \times \Omega(k^{1/99})$ grid minor, our lower bound can be extended to graphs of large treewidth as well. Let $\varphi^{\mathrm{tw}}(n, k, \beta(n))$ be the parametric complexity restricted to n-vertex graphs of treewidth k.

Corollary 4 (Graphs of large treewidth)**.** $\varphi^{\mathrm{tw}}(n, k, O((\log n)^3)) = k^{\Omega(\log k)}$.

Remark. In recent joint work with Pranabendu Misra, we show that $\varphi^{\mathrm{tw}}(n, k, \infty) = n^{O(k \log k)}$. This shows a dichotomy between the parametric complexity of graphs of small treewidth versus graphs of large treewidth. That is, graphs of constant treewidth have polynomial parametric complexity, and graphs of treewidth $\exp(\omega(\sqrt{\log n}))$ have superpolynomial parametric complexity.

To extend this result further, we require some candidate graphs when the treewidth is superconstant but bounded by $\exp(\omega(\sqrt{\log n}))$. The Mulmuley-Shah graphs [MS01] have small pathwidth but large parametric complexity, which make them good candidates. The following corollary is revealed simply by observing their graphs closely.

Corollary 5 (Corollary of [MS01]). *There is a graph on n vertices of pathwidth k having parametric complexity $n^{\Omega(\log k - \log\log n)}$.*

1.2 Significance of the main result

From their result, $\varphi(n, O((\log n)^3)) = n^{\Omega(\log n)}$, Mulmuley & Shah [MS01] 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 Appendix A for a discussion on this derivation).

Theorem 6. *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.*

Mulmuley & Shah observed that their result for the shortest path problem yields the same lower bound for the Weighted Bipartite Matching Problem [MS01, Corollary 1.1]. Our result extends this observation to planar graphs. 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 [Kas67] and Csanky [Csa75], and its remarkable recent application by Anari & Vazirani [AV18] (see also [San18]) to find perfect matchings in planar graphs. It is interesting that the lower bound for the Weighted Bipartite Matching Problem derived by Mulmuley & Shah continues to hold even when the input is restricted to be planar.

Parmetric 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 [NKBM06] 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 [CHKM17] 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.

- Erickson [Eri10] reformulates an $O(n \log n)$ time algorithm of Borradaile & Klein [BK09] for max-flows in planar graphs by considering parametric shortest path trees (see Karp & Orlin [KO81]) 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 a similar efficient algorithm because the tree might undergo $\Omega(n^2)$ changes.[1]

- Chakraborty, Fischer, Lachish & Yuster [CFLY10] 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.

Remark 1. Our construction yields a planar graph where 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 is also $n^{\Omega(\log n)}$.

Remark 2. 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 break points. Thus our result holds for *undirected planar graphs* as well.

1.3 Attacking Nikolova's conjecture

Before we move on to the proofs of our results, it is instructive to examine why our approach succeeded where earlier attempts failed.

Previous approaches: We recall two earlier efforts aimed towards resolving Nikolova's conjecture (Conjecture 2). In her PhD thesis, Nikolova [Nik09] considers embeddings of the planar graph in a plane, and shows that the edges can always be assigned weights in such a way that the number of break points is at least the number of faces in the embedding. Note, however, that the number of break points 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.* [CHKM17] observe that for series parallel graphs, Nikolova's conjecture is true; the parametric complexity of series-parallel graphs is in fact $O(n)$.

Our approach: It is instructive[2] to 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 $(1, 2, \ldots, n)$ in each layer. These arguments are based

[1] This is simply a closely related problem; our result has no implication on this.

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

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 $(p_1, p_2, \ldots, p_T)$ 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 [Nik09]). 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 7 (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 $(p_1, p_2, \ldots, p_T)$, where $i < j < k$. 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 the longest alternation-free sequence of paths in an n-vertex graph is an upper bound on $\varphi^{\mathrm{pl}}(n, \infty)$.

Definition 8. *Let $f(n, m)$ be the length of the longest alternation-free sequence paths in the layered graph $G[n, m]$; let $f^{\mathrm{pl}}(n, m)$ be the length of the longest alternation-free sequence of paths in any planar subgraph of $G[n, m]$ (with vertices s and t included).* ◊

Using the alternation-free property, one observes 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 [Car83b] and Mulumuley & Shah [MS01] imply that $f(n, n) \geq n^{\delta \log n}$ (for some $\delta > 0$). In Subsection 2.2, 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.*

Theorem 9. *(i) $f(n, 2^k) \geq n^k$, (ii) $f^{\mathrm{pl}}(n, (n-1)2^k) \geq n^k$.*

In Subsection 2.3, we present the construction for planar graphs in detail. This shows that the alternation-free property itself is insufficient to obtain significantly better upper bounds on $\varphi^{\mathrm{pl}}(n, \infty)$. While this construction provides some evidence against Nikolova's conjecture, it does not immediately refute it. There exist examples of alternation-free sequences of paths in planar graphs that do not arise as parametric shortest paths. Kuchlbauer [Kuc18, 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 [MS01]. 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 Subsection 3.1 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 break points 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 break points.

1.4 Updates from the previous version of this paper

This version of the paper differs from the version posted earlier in the following respects. (i) Some minor errors have been corrected. (ii) In the earlier version, the bound of $O((\log n)^3)$ on the bit lengths of the coefficients of the edge weights was argued informally; this version presents a formal argument. Furthermore, the coefficients are now integers, whereas earlier they were rational numbers. (iii) We attempt to provide a treewidth characterization for parametric shortest paths, which was not present in the earlier version. (iv) The appendix contains a proof that the parametric shortest path complexity of a $3 \times n$ grid graph is linear in n.

2 Alternation-free sequences of paths in graphs

In this section, we show Theorem 9. 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 [Car83b, MS01]. 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 2.1, we will formally build this connection between paths and words. In Subsection 2.2 and Subsection 2.3, we will use this connection to show Theorem 9, which implies the following.

Corollary 10 (Corollaries of Theorem 9)**.** *(i)* $f(n,n) \geq n^{\log n}$, *(ii)* $f^{\mathrm{pl}}(n, n^2 - n) \geq n^{\log n}$.

Viewing the construction behind the proof of Theorem 9 this from the treewidth lens, we have the following connection between pathwidth and alternation-free sequences in certain graphs.

Corollary 11 (Corollary of Theorem 9)**.** *There is an a graph on $n\kappa$ vertices of pathwidth κ for which there exists an alternation-free sequence of length $n^{\Omega(\log \kappa)}$.*

2.1 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

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

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

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

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 $\mathbb{Z}_n = \{0,1,\ldots,n-1\}$, 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 $G^{\text{npl}}[n,m]$:

$$E(G^{\text{npl}}) = (\{u\} \times L_0) \cup (L_0 \times L_1) \cup (L_1 \times L_2) \cup \cdots \cup (L_{m-2} \times L_{m-1}) \cup (L_{m-1} \times \{v\}).$$

Thus, for $j = 0,1,\ldots,m-2$, 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 $G^{\text{pl}}[n,m]$:

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

One can imagine that G^{pl} is drawn on the surface of a cylinder instead of the surface of a plane (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}). In G^{npl}, we may encode s-t paths by words in $\mathbb{Z}_n^m$: the word $\sigma = (\sigma_0, \sigma_1, \ldots, \sigma_{m-1}) \in \mathbb{Z}_n^m$ corresponds to the path

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

where $i_0 = \sigma_0$, and $i_{j+1} = i_j + \sigma_{j+1} \bmod n$, for $j = 0,1,\ldots,m-2$. 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 12 (Word). *Let $\mathbb{Z}_n$ denote the set $\{0,1,2,\ldots,n-1\}$ 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 $i \in \{0,1,\ldots,m-1\}$, let $\sigma[i]$ denote the i-th element of σ; let $\sigma[i:j]$ denote the subword $(\sigma[i], \sigma[i+1], \ldots, \sigma[j-1])$. 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* $i = 0, 1, \ldots, m-1$. *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* $\sigma \downarrow 3 = (7\ 3\ 2\ 3\ 6\ 3\ 2\ 3)$*.* ♢

Definition 13 (Alternation-free sequence of words)**.** *Let* S *be sequence of* ℓ *words from* $\mathbb{Z}_n^m$*, that is,* $S \in (\mathbb{Z}_n^m)^\ell$*. 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 any such alternation we must have either $v = m$ *or* $v - u \geq 2$*. If* S *has no alternation, we say it is* alternation-free*.* ♢

Proposition 14 (Paths from sequences)**.** *If* $S = (\sigma_i : i = 0, 1, \ldots, \ell - 1) \in (\mathbb{Z}_n^m)^\ell$ *is an alternation-free sequence of words, then* $(p_{\sigma_i} : i = 0, 1, \ldots, \ell - 1)$ *is an alternation-free sequence of paths in* G_m^{npl}*. Similarly, if* $T = (\tau_i : i = 0, 1, \ldots, \ell - 1) \in (\{0, 1\}^m)^\ell$ *is an alternation-free sequence of words, then* $(p_{\tau_i} : i = 0, 1, \ldots, \ell - 1)$ *is an alternation-free sequence of paths in* G_m^{pl}*.*

Proof. Straightforward. Note that the case $v = m$ in the second condition of Definition 13 is used to verify that there is no alternation involving pairs of vertices of the form (u, t). □

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

2.2 Construction of 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 $\{0, 1\}$ respectively. We first describe X. The i-th word ($i = 0, 1, \ldots, n^\ell - 1$) of X is given by $X[i] = (b_0) \downarrow b_1 \downarrow \cdots \downarrow b_{\ell-1}$, where $(i)_n = \sum_{j=0}^{\ell-1} b_j n^j$ is the base n representation of i. For example, suppose $n = 4$ and $i = 114$. Then $X[114] = (2) \downarrow 0 \downarrow 3 \downarrow 1 = (2\ 1\ 3\ 1\ 0\ 1\ 3\ 1)$ because 114 is equal to 1302 in base 4.

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 $\hat{\imath} = 1^i 0^{n-1-i} \in \{0, 1\}^{n-1}$. 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 $\hat{X}[114] = (110\ 100\ 111\ 100\ 000\ 100\ 111\ 100)$. Now we will show that X and $\hat{X}$ are alternation-free.

Lemma 15. *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*) $T = (S \downarrow 0) \circ (S \downarrow 1) \circ \cdots \circ (S \downarrow (n-1))$ *is an alternation-free sequence of* $n\ell$ *words, where each word is in* $\mathbb{Z}_n^{2m}$*.*[3]

[3]Here $\circ$ represents concatenation of sequences.

Proof. For part (a), note that if $S \downarrow j$ has an alternation at (a, b, c) ($0 \le a < b < c \le \ell - 1$) between (s, t) ($0 \le s < t \le 2m$), then S itself has an alternation at (a, b, c), between $(\lceil s/2 \rceil, \lceil t/2 \rceil)$. 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) ($0 \le a < b < c \le n\ell$) between (s, t) ($0 \le s < t \le 2m$). If σ_a and σ_c have the same symbol in their odd positions 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 (s, t).

On the other hand, suppose σ_a and σ_c have different symbols in their odd positions. Since $\sigma_a[t-1] = \sigma_c[t-1]$, we conclude that t is odd. In particular, $t \neq 2m$ and thus $t - s \ge 2$ (as observed at the end of Definition 13). This means that the interval $\{s, s+1, \ldots, t-1\}$ includes an odd number. Hence $\sigma_a[s : t] \neq \sigma_c[s : t]$, and there is no alternation at (a, b, c) between (s, t). □

Theorem 16. *For all $\ell \ge 1$, there is an alternation-free sequence T of n^ℓ words in $\mathbb{Z}_n^{2^\ell}$.*

Proof. We will use Lemma 15 and induction on ℓ. For $\ell = 1$, the alternation-free sequence is simply $T = (0, 1, 2, \ldots, n-1)$, 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

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

By Lemma 15, T is an alternation-free sequence of $n \cdot n^\ell = n^{\ell+1}$ words in $\mathbb{Z}_n^{2 \cdot 2^{\ell-1}} = \mathbb{Z}_n^{2^\ell}$. □

Note that the sequence T is exactly the sequence X that we described above.

2.3 Construction of alternation-free sequences of binary words

Consider the following *unary* encoding, where we map $i \in \mathbb{Z}_n$ to the binary word $\hat{\imath} = 1^i 0^{n-1-i} \in \{0, 1\}^{n-1}$. Let $\hat{\mathbb{Z}}_n = \{\hat{0}, \hat{1}, \ldots, \widehat{n-1}\}$. 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 17. *Suppose $\hat{S} \in (\hat{\mathbb{Z}}_n^m)^\ell$ is alternation-free sequence of ℓ binary words of length $m(n-1)$ each. Then,*

(a) *for all $\hat{\jmath} \in \hat{\mathbb{Z}}_n$, $\hat{S} \downarrow \hat{\jmath}$ is an alternation-free sequence of ℓ words in $\{0, 1\}^{2m(n-1)}$;*

(b) *$\hat{T} = (\hat{S} \downarrow \hat{0}) \circ (\hat{S} \downarrow \hat{1}) \circ \cdots \circ (\hat{S} \downarrow \widehat{n-1})$ is an alternation-free sequence of $n\ell$ words in $\{0, 1\}^{2m(n-1)}$.*

Proof. For part (a), consider $\hat{S} \downarrow \hat{\jmath}$. 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, the odd numbered blocks all contain the word $\hat{\jmath}$. 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 $\hat{S} \downarrow \hat{\jmath}$ is alternation-free because $\hat{S}$ is known to be alternation-free. We now make this idea more precise. Suppose $\hat{S} \downarrow \hat{\jmath} = (\sigma_i : i = 0, 1, \ldots, \ell - 1)$ has an alternation at (a, b, c) $(0 \le a < b < c \le \ell - 1)$ between (s, t) $(0 \le s < t \le 2m(n-1))$. Suppose $t = 2m(n-1)$, that is, it points to the end of the word. Then s cannot be a location in the last block, for the entire block is identical in all words in $\hat{S} \downarrow \hat{\jmath}$. Suppose $s = q(n-1) + r$, where $r = s \bmod n - 1$ and $q < 2m - 1$. We conclude that S has an alternation at (a, b, c) between $(\lceil q/2 \rceil (n-1) + r, m(n-1))$, contradicting our assumption that S is alternation-free. So we may assume that $t < 2m(n-1)$. We may also assume that (s, t) has been chosen so that $t - s$ is minimal. This implies that $\sigma_a[s] = \sigma_c[s] \ne \sigma_b[s]$, and similarly that $\sigma_a[t-1] = \sigma_c[t-1] \ne \sigma_b[t-1]$. In particular, both s and $t-1$ are indices into even numbered blocks. Suppose $s = q(n-1) + r$, $t = q'(n-1) + r'$, where $r = s \bmod n - 1$ and $r' = t \bmod n - 1$. 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 our assumption that $\hat{S}$ is alternation-free. This establishes part (a).

For part (b), suppose $\hat{T}$ has an alternation at (a, b, c) $(0 \le a < b < c \le n\ell - 1)$ between (s, t) $(0 \le s < t \le 2m(n-1))$; assume that (s, t) has been chosen so that $t - s$ is minimal. Now $\hat{T}$ consists of subsequences of words of the form $\hat{S} \downarrow \hat{\jmath}$. We have two cases. First, suppose σ_a, σ_b and σ_c come from the same subsequence of the form $\hat{S} \downarrow \hat{\jmath}$. Then, part (a) gives us the necessary contradiction. So we may assume that the odd numbered blocks of σ_a and σ_c are different. Note that the contents of the odd numbered blocks are monotonically non-decreasing in $\hat{\mathbb{Z}}_n$, so if there is a full odd numbered block in the range $\{s, s+1, \ldots, t-1\}$, then $\sigma_a[s:t] \ne \sigma_c[s:t]$, and there can be no alternation. So we may assume that there is no full odd numbered in the range $\{s, s+1, \ldots, t-1\}$. To complete the proof, we now use two crucial facts about our encoding. First, in the sequence of words $(\hat{0}, \hat{1}, \ldots, \widehat{n-1})$, the bit at any position starts at 0 and flips to 1, never to return to 0 again. This implies that, in fact, both s and $t-1$ lie within the same even numbered block. Second, our encoding has the property that if for two words $\hat{\jmath}_1, \hat{\jmath}_2 \in \hat{\mathbb{Z}}_n$, we have $|\hat{\jmath}_1[u:v]|_1 = |\hat{\jmath}_2[u:v]|_1$, then $\hat{\jmath}_1[u:v] = \hat{\jmath}_2[u:v]$. This implies that $\sigma_a[s:t] = \sigma_b[s:t] = \sigma_c[s:t]$; so there is no alternation. This establishes part (b). □

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

3 The planar construction with linearly varying edge weights

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 [MS01], which in turn was based on the construction of Carstensen [Car83b]. 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 break points in the lower envelope by a factor n. After m steps of induction, we obtain a graph with $\mathsf{poly}(n) \cdot \exp(O(m))$ vertices and n^m break points. Figure 3 illustrates this assembly for $m = 3, n = 3$, following the template of Figure 1. We will explain this construction in detail later. The edge weights in the constituent graphs in Figure 1 are carefully chosen, but are not

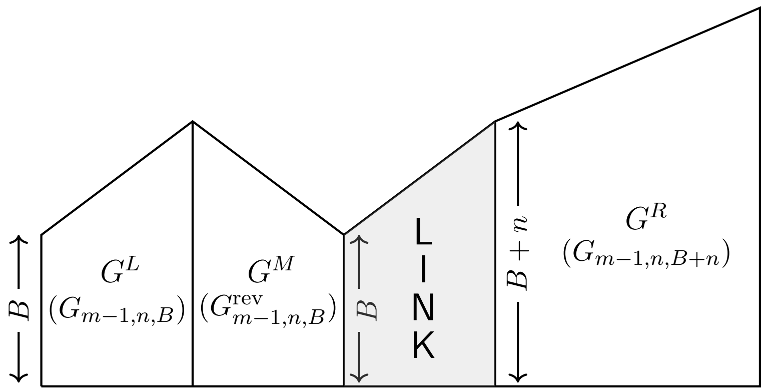

Figure 1: $G_{m,n,B}$ is obtained by composing $G_{m-1,n,B}$, $G^{\text{rev}}_{m-1,n,B}$, the linking gadget, and $G_{m-1,n,B+n}$

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 4, we show that the new edge weights in LINK also result in a large number of break points. In Subsection 3.1, we show that the non-planar graph G^{npl} 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.

3.1 The linking gadget

A linking gadget $\mathbf{L}(B, n)$ is a bipartite graph $G(U, V, E, (w_e : e \in E))$ with $U = \{0, 1, \ldots, B-1\}$, $V = \{0, 1, \ldots, B+n-1\}$, $E = \{(b, b+r) : b \in U, r = 0, 1, \ldots, n\}$. 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 2). 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} of G. 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

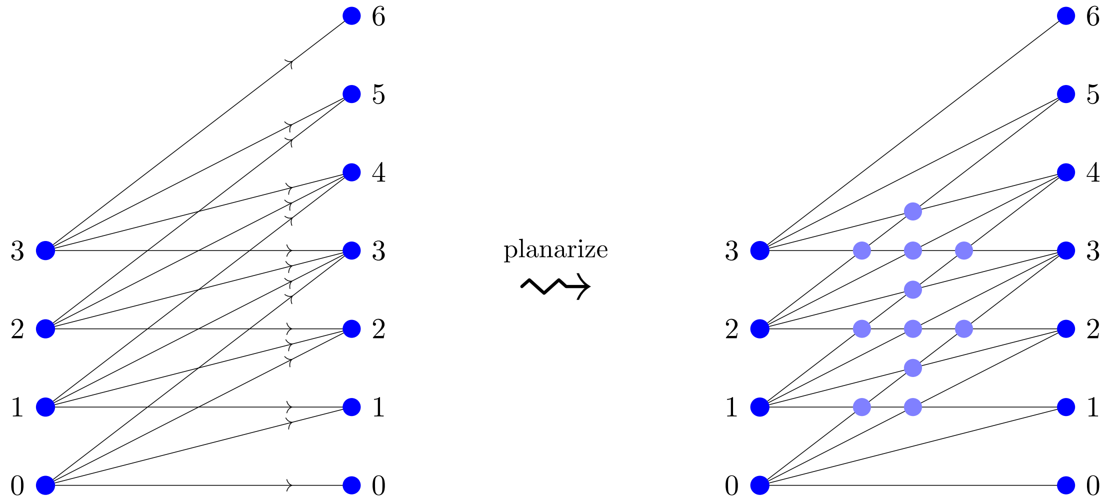

Figure 2: The LINK *gadget* $\mathbf{L}(B, n)$ *for* $B = 4, n = 3$ *and its planarization* $\mathbf{L}(B, n)$

most n, which leads to the following observation.

Fact 18. *The 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 21). Before that, we need to formalize what it means for G^{pl} to mimic G.

Definition 19. *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*.* ◊

Lemma 20. *Suppose* $J : E(G) \to \mathbb{Z}$*, 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*

$$w_{b,b+r} = J(b, b + r) + K \left(\frac{r(r + 1)}{2}\right) + Lr\lambda, \qquad \text{where } 0 \leq b \leq B - 1 \text{ and } 0 \leq r \leq n.$$

Then G^{pl} *faithfully simulates* G*.*

Proof. Consider vertices $b \in U$ and $j \in V$ such that $0 \le b \le j \le b+n$. 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 21. $\mathcal{C}(Q) - \mathcal{C}(P) \ge n^{-2}K - 2\max_{e\in E(G)} |J(e)|$.

Proof of Claim 21. Let Q consist of vertices $q_0 = (x_0,y_0), q_1 = (x_1,y_1), q_2 = (x_2,y_2), \ldots, q_t = (x_t,y_t)$, where $(x_0,y_0) = (0,b)$ and $(x_t,y_t) = (1,j)$ for some b,j. For $i = 1,2,\ldots,t$, 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 $i \in 1,2,\ldots,t$, we have

$$\rho_i \ge n^{-2}; \qquad \text{(using Fact 18)}$$

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

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

$$w_{q_{i-1},q_i} = \left(J(y, y+r_i) + K\left(\frac{r_i(r_i+1)}{2}\right) + Lr_i\lambda\right)(x_i - x_{i-1}) \qquad (\text{where } y = y_i - r_i x_i).$$

Since $0 \le \rho_i \le 1$ and $\sum_{i=1}^{t} \rho_i = 1$, we may define a random variable $\mathbf{i}$, that takes the value $i \in \{0,1,\ldots,n\}$ with probability ρ_i.

$$\mathcal{C}(Q) - \mathcal{C}(P) \ge -2\max_e |J(e)| + K\left(\mathbb{E}\left[\frac{r_{\mathbf{i}}^2}{2}\right] - \frac{r^2}{2}\right) + \left(\frac{K+2L\lambda}{2}\right)(\mathbb{E}[r_{\mathbf{i}}] - r) \tag{22}$$

$$\ge -2\max_e |J(e)| + K\left(\mathbb{E}\left[\frac{r_{\mathbf{i}}^2}{2}\right] - \frac{r^2}{2}\right) \qquad (\text{because } \mathbb{E}[r_{\mathbf{i}}] = r) \tag{23}$$

$$\ge -2\max_e |J(e)| + \frac{K}{2}\mathbf{var}[r_{\mathbf{i}}]. \tag{24}$$

We show a lower bound for $\mathbf{var}[r_{\mathbf{i}}]$. 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,

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

Combining this with (24) establishes Claim 21. □

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 20. □

4 Proof of the main result

We now prove Theorem 1. The proof is by induction. We begin with a graph on n vertices with n disjoint intervals. In each inductive step, we roughly triple the size of the graph and subdivide each interval into n intervals. After $\log n$ steps, we end up with a graph on $\mathsf{poly}(n)$ vertices with $n^{\log n}$ intervals. For simplicity, we assume that $n \geq 4$. Throughout our proof, we work with a fixed value of n.

4.1 Inductive definition of intervals

In our recursive construction, we will construct paths that reign as the shortest path in particular intervals (that is, 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 shortest path in each interval. This will establish that there are many break points in the cost of the shortest path in our graph. The graph we construct and the role of the intervals is described in detail below. We first place the framework by describing the intervals inductively. The intervals depend on a parameter N, defined as

$$N = n^2. \tag{25}$$

Then, for $m = 0, 1, \ldots, \lfloor \log n \rfloor$ and $j = 0, 1, \ldots, n^m - 1$, 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 [MS01] 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 $(\mathcal{I}(j,m) : j = 0, 1, \ldots, n^m - 1)$ 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, $\alpha(j,m) \leq N^{m+1}$.

4.2 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 $1, n+1, 2n+1, 3n+1, \ldots$, and $b \in \{0, 1, 2, \ldots, B-1\}$ 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 the vertices in the last layer of the graph to a new vertex t using edges of weight 0, so that we have pristine s-t paths as promised. Thus, we will 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). m helps ensure that the number of break points in the cost of the shortest path is large.

See Figure 3 for a step-by-step visualization of this construction. The formal induction will be carried out using a predicate Φ, which we now define.

<u>The Predicate Φ</u>

For B, D and 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 paths P_{bj} (for $b = 0, 1, \ldots, B-1$ and $j = 0, 1, \ldots, n^m - 1$) 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

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

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

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

(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 in the list $(P_{bj} : b = 0, 1, \ldots, B-1)$ are vertex-disjoint.

(vi) For all b, the paths in the list $(P_{bj} : j = 0, 1, \ldots, n^m - 1)$ are distinct.

The following lemma is essentially the same as Lemma 4.1 of Mulmuley & Shah [MS01]. 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 Subsection 3.1 and Lemma 20 to ensure that our graphs are planar.

Lemma 26 (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 26). Using property (i), the number of vertices in the corresponding graph $G(1, 0, \lfloor \log n \rfloor)$ is at most

$$\begin{aligned}
(3^{m+1} - 1)(B + mn)^4 &\leq (3^{\log n + 1} - 1)(1 + (\log n)n)^4 \\
&\leq (3n^{1.585})(1 + (\log n)n)^4 && \text{(since } \log_2 3 \leq 1.585) \\
&\leq 6n^{1.585}(n \log n)^4 && \text{(since } n \geq 4) \\
&\leq 6n^{1.585}(n \cdot n^{0.6})^4 \\
&\leq 6n^8.
\end{aligned}$$

To this graph we attach a sink vertex t as stated above. 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 $n^{\lfloor \log n \rfloor}$ break points.

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

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

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

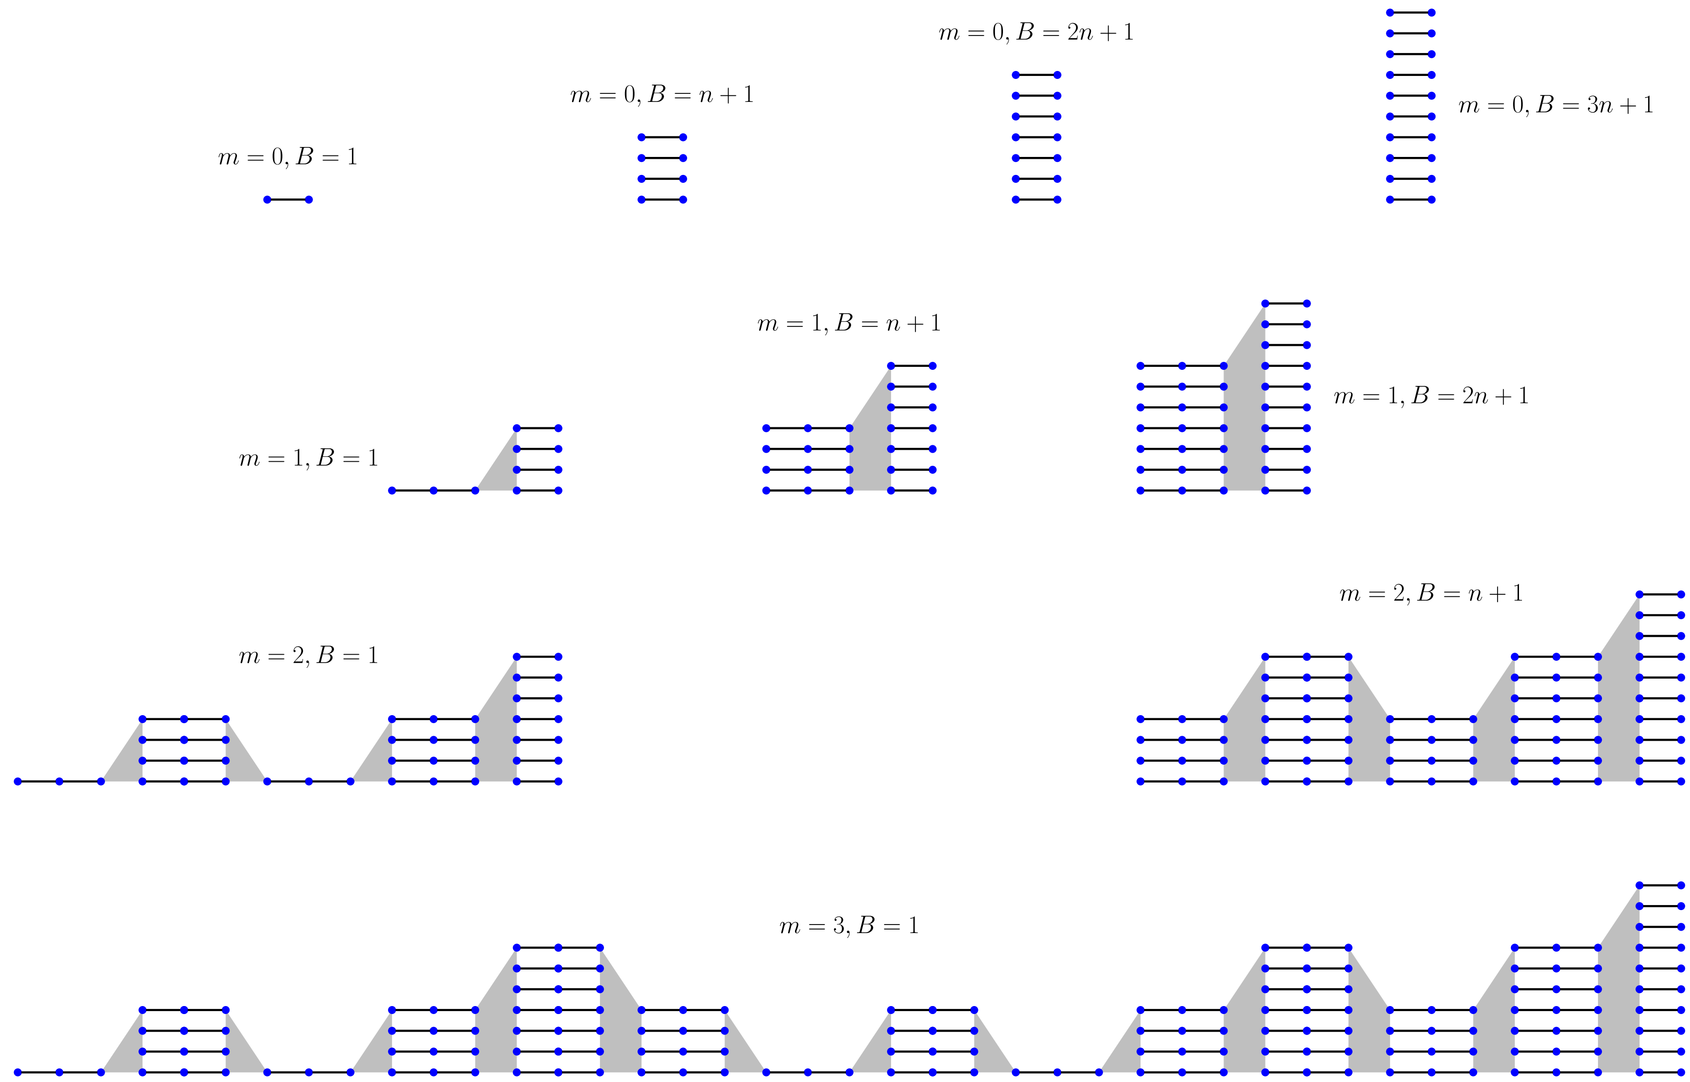

Figure 3: An instantiation of the graph $G_{m,n,B}$ for $m = 3, n = 3, B = 1$

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 $\nu^{\Omega(\log \nu)}$ break points. 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. □

4.3 Proof of the main lemma

Proof of Lemma 26. 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 $m \geq 1$ and 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

$$K_L = 400\, N^{m+4} B^2; \tag{27}$$

$$K_R = 20\, N^3 B; \tag{28}$$

$$D_L = \frac{N}{2K_L}\left(D - \frac{K_R}{N}\right); \tag{29}$$

$$D_R = 1. \tag{30}$$

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

Construction of G: The graph G is built by 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^L_{bj} where $0 \leq b < B$ and $0 \leq j < n^{m-1}$. 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 1). 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^R_{bj} where $0 \le b < B+n$ and $0 \le j < n^{m-1}$.

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^L_e, w^M_e and w^R_e denote the weights of the edges in G^L, G^M and G^R, and let w_e denote their weights in G.

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

In essence, we are *scaling* (by factors K_L and K_R) and *stretching* (by a factor N)[4] 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

$$w_{b,b+r} = NDrb + \frac{K_R}{N}\left(\left(\frac{r(r+1)}{2}\right)N - r\lambda\right), \qquad \text{where } 0 \le b \le B-1 \text{ and } 0 \le r \le n, \tag{31}$$

and let $\mathbf{L}^{\mathrm{pl}}(B,n)$ be its planarized version. Note that the D used in (31) 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, $\mathbf{L}^{\mathrm{pl}}$ and G^R is shown in Figure 1. Since G^L, G^M and G^R are planar by induction, and $\mathbf{L}^{\mathrm{pl}}(B,n)$ is planar, the graph obtained by composing them is also planar (Figure 3). This completes the description of all the constituent components of G.

Before we proceed further, let us verify that for our choice of parameters, $\mathbf{L}^{\mathrm{pl}}$ faithfully simulates its non-planar counterpart. Invoke Lemma 20 with $J(b,b+r) = NDrb$, $K = K_R$ and $L = -K_R/N$. For the setting of K_R in (28), we have (recall $N = n^2$)

$$n^2\left(1 + 2\max_{e\in E(\mathbf{L}^{\mathrm{pl}}(B,n))}|J(e)|\right) \le n^2\,(1 + 2N|D|nB) \le 4(|D|+1)N^{2.5}B \le K_R,$$

so the requirement $K \ge n^2\,(1 + 2\max_e |J(e)|)$ of Lemma 20 holds. Thus, $\mathbf{L}^{\mathrm{pl}}(B,n)$ faithfully simulates $\mathbf{L}(B,n)$.

Finally, in order invoke the induction hypothesis for $\Phi(B, D_L, m-1)$ and $\Phi(B+n, D_R, m-1)$, we need to show that $|D_L| \le 1$ ($|D_R| = 1$ from (30)).

$$\begin{aligned}
|D_L| &= \left|\frac{N}{2K_L}(D - 20N^2B)\right| && \text{(from (29))}\\
&\le \left|\frac{ND}{800N^{m+4}B^2}\right| + \left|\frac{20N^3B}{800N^{m+4}B^2}\right| && \text{(from (27))}\\
&\le \frac{1}{800} + \frac{1}{40} \ll 1. && \text{(since } |D| \le 1)
\end{aligned}$$

[4] Recall that our intervals are stretched by a factor N when we go from one level of recursion to the next (Subsection 4.1).

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 1)

$$G = G^L \circ G^M \circ \mathbf{L}^{\rm pl} \circ G^R, \tag{32}$$

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 $0 \le j < n^m$, write $j = nd + r$ with $0 \le d < n^{m-1}$ and $0 \le r < n$; then for $0 \le b < B$, let

$$P_{bj} = P^L_{bd} \circ (P^L_{bd})^{\rm rev} \circ \texttt{link}(b, b+r+1) \circ P^R_{b+r+1,d}, \tag{33}$$

where $\texttt{link}(b, b+r+1)$ is the unique shortest path (the straight line) in $\mathbf{L}^{\rm pl}$ 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 (33) 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 $\mathbf{L}^{\rm pl}(B,n)$ adds at most $(B+n)^4$ new vertices. This means that the number of vertices in the planarized version of G is at most

$$\underbrace{2(3^m-1)(B+(m-1)n)^4}_{G^L, G^M} + \underbrace{(B+n)^4}_{\mathbf{L}^{\rm pl}(B,n)} + \underbrace{(3^m-1)(B+n+(m-1)n)^4}_{G^R} \le (3^{m+1}-1)(B+mn)^4.$$

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,

$$\max_e\{|\mathsf{num}(a_{1e})|, |\mathsf{num}(a_{2e})|, |\mathsf{num}(b_{1e})|, |\mathsf{num}(b_{2e})|\} \le (400\,NB')^{5(m-1)^2},$$
$$\max_e\{|\mathsf{den}(a_{1e})|, |\mathsf{den}(a_{2e})|, |\mathsf{den}(b_{1e})|, |\mathsf{den}(b_{2e})|\} \le 2^{m-1}n^2,$$

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, $\mathbf{L}^{\rm pl}(B,n)$ 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,

$$a^L_e = a^L_{1e} + a^L_{2e}\left(\frac{N}{2K_L}\left(D - \frac{K_R}{N}\right)\right). \qquad \text{(before scaling)}$$
$$b^L_e = b^L_{1e} + b^L_{2e}\left(\frac{N}{2K_L}\left(D - \frac{K_R}{N}\right)\right). \qquad \text{(before scaling)}$$

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.

$$\begin{aligned}
a_e &= K_L a_{1e}^L + \frac{a_{2e}^L N}{2}\left(D - \frac{K_R}{N}\right) && \text{(after scaling)}\\
&= 400\, N^{m+4} B^2 a_{1e}^L + \frac{a_{2e}^L N}{2}\left(D - \frac{20\, N^3 B}{N}\right) && \text{(using (27))}\\
&= \underbrace{\left(400\, N^{m+4} B^2 a_{1e}^L - 10\, a_{2e}^L N^3 B\right)}_{a_{1e}} + \underbrace{\left(\frac{a_{2e}^L N}{2}\right)}_{a_{2e}} D.
\end{aligned}$$

$$\begin{aligned}
b_e &= \frac{K_L b_{1e}^L}{N} + \frac{b_{2e}^L}{2}\left(D - \frac{K_R}{N}\right) && \text{(after scaling)}\\
&= \frac{400\, N^{m+4} B^2 b_{1e}^L}{N} + \frac{b_{2e}^L}{2}\left(D - \frac{20\, N^3 B}{N}\right) && \text{(using (27))}\\
&= \underbrace{\left(400\, N^{m+3} B^2 b_{1e}^L - 10\, b_{2e}^L N^2 B\right)}_{b_{1e}} + \underbrace{\left(\frac{b_{2e}^L}{2}\right)}_{b_{2e}} D.
\end{aligned}$$

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.

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

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 $(b, b+r) \in E(\mathbf{L}(B,n))$.

$$\begin{aligned}
w_{b,b+r} &= NDrb + \frac{K_R}{N}\left(\binom{r(r+1)}{2} N - r\lambda\right) && \text{(using (31))}\\
&= NDrb + \frac{r(r+1)K_R}{2} - \frac{rK_R}{N}\lambda\\
&= \underbrace{\left(NDrb + 10\, r(r+1)N^3 B\right)}_{a_e} + \underbrace{\left(-20\, rN^2 B\right)}_{b_e}\lambda.
\end{aligned}$$

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 18). As for the numerator,

$$\begin{aligned}
|NDrb + 10\, r(r+1)N^3 B| &\le NDnB + 10N^4 B && \text{(since } 0 \le b \le B-1, 0 \le r \le n, N = n^2)\\
&\le (400NB)^{5m^2}.
\end{aligned}$$

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.

$$
\begin{aligned}
a_e^R &= a_{1e}^R + a_{2e}^R D_R. && \text{(before scaling)} \\
b_e^R &= b_{1e}^R + b_{2e}^R D_R. && \text{(before scaling)} \\
a_e &= K_R a_{1e}^R + K_R a_{2e}^R && \text{(after scaling)} \\
&= \underbrace{20\,N^3 B a_{1e}^R + 20\,N^3 B a_{2e}^R}_{a_{1e}}. && \text{(using (28))} \\
b_e &= \frac{K_R}{N} b_{1e}^R + \frac{K_R}{N} b_{2e}^R && \text{(after scaling)} \\
&= \underbrace{20\,N^2 B a_{1e}^R + 20\,N^2 B a_{2e}^R}_{b_{1e}}. && \text{(using (28))}
\end{aligned}
$$

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.

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

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

Properties (v), (vi): Given our definition of P_{bj} (33), properties (v) and (vi) are straightforward to verify. We now verify property (iv).

Property (iv): In the subsequent calculations, we use the following notation. Paths of G are composed of paths coming 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 of N. Thus, the cost of the P_{bj} in the graph G is given by

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

Substitute $b = 0$ to get

$$\mathcal{C}(P_{0j})(\lambda) = 2K_L\mathcal{C}_L(P^L_{0d})\left(\frac{\lambda}{N}\right) + \frac{K_R}{N}\left[\frac{(r+1)(r+2)}{2}N - (r+1)\lambda\right] + K_R\mathcal{C}_R(P^R_{0d})\left(\frac{\lambda}{N}\right) + K_R(r+1)D_R\alpha(d,m-1).$$

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

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

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 34 and Claim 35).

In Claim 34, we track paths from an input vertex as they travel through G^L and G^M. In Claim 35, we analyze how such paths continue through $\mathbf{L}^{\rm pl}$ and G^R. Fix a j ($0 \le j \le n^m - 1$, say $j = nd + r$, for $0 \le d \le n^{m-1} - 1$ and $0 \le r \le n-1$) 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 34. *Let Q be a path from the input vertex b to the last layer of $G^L \circ G^M$ (note that $P^L_{bd} \circ (P^L_{bd})^{\rm rev}$ is one such path). If $Q \neq P^L_{bd} \circ (P^L_{bd})^{\rm rev}$, then*

$$\mathcal{C}(Q) - \mathcal{C}(P^L_{bd} \circ (P^L_{bd})^{\rm rev}) \ge K_L/2.$$

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,

$$\begin{aligned}
\mathcal{C}(Q) - \mathcal{C}(P^L_{bd} \circ (P^L_{bd})^{\rm rev}) &= \left(\mathcal{C}(Q^L) + \mathcal{C}(Q^M)\right) - \left(\mathcal{C}(P^L_{bd}) + \mathcal{C}((P^L_{bd})^{\rm rev})\right) \\
&= \left(\mathcal{C}(Q^L) - \mathcal{C}(P^L_{bd})\right) + \left(\mathcal{C}(Q^M) - \mathcal{C}(P^L_{cd})\right) + \left(\mathcal{C}((P^L_{cd})^{\rm rev}) - \mathcal{C}((P^L_{bd})^{\rm rev})\right) \\
&\ge \overbrace{\mathcal{C}(Q^L) - \mathcal{C}(P^L_{bd})}^{\text{Term I}} + \overbrace{\mathcal{C}(Q^M) - \mathcal{C}(P^L_{cd})}^{\text{Term II}} - \overbrace{K_LBD_L\alpha(d,m-1)}^{\text{Term III}}.
\end{aligned}$$

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 $Q \neq P^L_{bd} \circ (P^L_{bd})^{\rm rev}$, then one of the following is true.

(a) $Q^L \neq P^L_{bd}$;

(b) $c = b$ and $Q^M \neq (P^L_{bd})^{\text{rev}}$;

(c) $c \neq b$ and $Q^M \neq (P^L_{cd})^{\text{rev}}$ (here we use the fact that the paths P^L_{bd} and P^L_{cd} 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 $\geq K_L$. Also note that $\alpha(d, m-1) \leq N^m$. For the setting of K_L according to (27) we have $|\text{Term III}| = |K_L B D_L \alpha(d, m-1)| \ll K_L/10$. This completes the proof of Claim 34. $\square$

Since K_L is positive, Claim 34 implies that $P^L_{bd} \circ (P^L_{bd})^{\text{rev}}$ 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 (33).

Claim 35. *Let $\lambda \in \mathcal{I}(j, m)$, where $j = nd + r$ ($0 \leq d < n^{m-1}$ and $0 < r \leq n-1$). Let P be a path from the input vertex b of $\mathbf{L}(B, n)$ to the last layer of G^R (note that $\texttt{link}(b, b+r+1) \circ P^R_{b+r+1,d}$ is one such path). If $P \neq \texttt{link}(b, b+r+1) \circ P^R_{b+r+1,d}$, then*

$$\mathcal{C}(P)(\lambda) - \mathcal{C}(\texttt{link}(b, b+r+1) \circ P^R_{b+r+1,d})(\lambda) \geq 1\,. \tag{36}$$

Proof of claim. Fix the input vertex b. The induction hypothesis guarantees that P^R_{xd} 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 travels along the shortest path in G^R, that is, it has the form

$$P_k = \texttt{link}(b, b+k) \circ P^R_{b+k,d},$$

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

$$Z_0 \gg Z_1 \gg \cdots \gg Z_r \gg Z_{r+1} \ll Z_{r+2} \ll \cdots \ll Z_n, \tag{37}$$

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 $k = 1, 2, \ldots, n$, we have

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

where

$$w_{b,b+k} = NDkb + \frac{K_R}{N}\left(\left(\frac{k(k+1)}{2}\right)N - k\lambda\right);$$
$$w_{b,b+k-1} = ND(k-1)b + \frac{K_R}{N}\left(\left(\frac{(k-1)k}{2}\right)N - (k-1)\lambda\right);$$
$$\mathcal{C}(P^R_{b+k,d})(\lambda) = K_R\left[\mathcal{C}_R(P^R_{0,d})(\lambda/N) + (b+k)D_R\alpha(d,m-1)\right];$$
$$\mathcal{C}(P^R_{b+k-1,d})(\lambda) = K_R\left[\mathcal{C}_R(P^R_{0,d})(\lambda/N) + (b+k-1)D_R\alpha(d,m-1)\right].$$

Thus,

$$Z_k - Z_{k-1} = NDb + \frac{K_R}{N}(kN-\lambda) + K_R D_R \alpha(d,m-1) \quad (38)$$
$$= NDb + \frac{K_R}{N}(kN + N\alpha(d,m-1) - \lambda) \qquad \text{(recall } D_R = 1) \quad (39)$$
$$= NDb + \frac{K_R}{N}(\alpha(k-1,m) - \lambda). \quad (40)$$

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

$$\alpha(k-1,m) - \lambda \in [\alpha(k-1,m) - \alpha(r,m) - N + 1, \alpha(k-1,m) - \alpha(r,m) - 1] \quad (41)$$
$$= [(k-(r+1))N - N + 1, (k-(r+1))N - 1]. \quad (42)$$

Thus, for $k = 1, 2, \ldots, r+1$, we have $\alpha(k-1,m) - \lambda \le -1$ and for $k = r+2, \ldots, n$, we have $\alpha(k-1,m) - \lambda \ge +1$. Returning to (40) with this, we obtain

$$Z_k - Z_{k-1} \le NDb - \frac{K_R}{N} \qquad \text{for } k = 1, \ldots, r+1, \text{ and} \quad (43)$$
$$Z_k - Z_{k-1} \ge NDb + \frac{K_R}{N} \qquad \text{for } k = r+2, \ldots, n. \quad (44)$$

Since $K_R \gg N^2 b$ and $-1 \le D \le +1$, the RHS of (43) is negative and the RHS of (44) is positive. This confirms (37) and establishes Claim 35. □

We are now in a position to establish property (iii) and complete the induction. By Claim 34, if the shortest path from b deviates from $P^L_{bd} \circ (P^R_{bd})^{\text{rev}}$ in $G^L \circ G^M$, then the increase in cost is at least $K_L/2$. We will 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 shortest path from the input vertex b in G when restricted to $G_L \circ G_M$ to be precisely $P^L_{bd} \circ (P^R_{bd})^{\text{rev}}$. 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,

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

Thus, the shortest path in G must follow the path $P^L_{bd} \circ (P^R_{bd})^{\text{rev}}$ 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 $\mathcal{C}(P_{bd})$. Claim 35 now confirms that it must continue by taking the edge $\texttt{link}(b, b+r+1)$ and $P^R_{b+r+1,d}$; 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 (33). This completes the proof of property (iii), hence completing the proof of Lemma 26. □

5 Upper bound for polynomially varying edge weights

Throughout this paper, the edge weights are linear functions of λ. In this section, we will show that even if the edge weights are allowed to be polynomials of degree d in the parameter λ, the upper bound is not significantly higher than that for $d = 1$. Let $\varphi_d(n, \beta(n))$ be the maximum possible number number of break points in $\mathcal{C}(\lambda)$ for an n-vertex graph when the edge weights are polynomials of degree at most d in a parameter λ, where the bit lengths of the coefficients in the weights of the edges are bounded by β.

Theorem 45. *For all fixed d, we have $\varphi_d(n, \infty) = n^{\log n + (\alpha(n) + O(1))^d}$, where $\alpha(n)$ is the extremely slow growing inverse Ackermann function.*

Proof. We adapt to our setting an argument due to Gusfield [Car83b, Page 100] (later, a similar argument was made by Dean [Nik09, Page 86]). Fix a $d > 1$ and consider only those graphs whose edge weights are polynomials of degree at most d. Let $f(n, m)$ be the maximum length of a sequence of shortest paths[5], when the paths are restricted to have at most m edges. Let $m = 2^k$, and 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 2^{k-1} edges from s to v, and at most 2^{k-1} edges from v to t. Then, the number of such paths in σ that pass through v is at most $2f(n, 2^{k-1})$. Accounting for all v, we obtain that there are at most $2nf(n, 2^{k-1})$ paths in the sequence σ. Note that σ might have alternations. Thus, if N is the number of *distinct* paths in σ, then $N \le 2nf(n, 2^{k-1})$. Since the costs of these paths are polynomials of degree at most d in λ, two paths can alternate at most $d + 1$ times (two distinct degree d polynomials

[5]Note 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).

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 [Mat02, 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 any n-vertex graph), we have $\alpha(N) \ll \alpha(n) + 2$. Thus, $f(n, 2^k) \leq N \cdot 2^{\alpha(N)^d} \leq 2nf(n, 2^{k-1}) \cdot 2^{(\alpha(n)+2)^d}$, which yields $f(n, 2^k) \leq (2n)^k \cdot f(n, 1) \cdot 2^{k(\alpha(n)+2)^d}$. Our theorem follows from this by substituting $f(n, 1) \leq 1$ and taking $k = \lceil \log n \rceil$. □

Note that our proof works for any family of functions $\mathcal{F}$ which satisfy the following conditions (the family of polynomials obviously 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.

6 Conclusion

In this work, we studied (among other things) how the parametric shortest path complexity changes with the topology of the graph. As we discovered, treewidth is the right measure. Graphs with small (constant) treewidth have small (polynomial) parametric complexity, and graph with large ($\exp(\omega(\sqrt{\log n}))$) treewidth have superpolynomial parametric complexity. However, an unexplored gap still remains. The following conjecture, in particular, is interesting.

Conjecture 46. *There is an n-vertex graph of treewidth k with parametric complexity $n^{\Omega(\log k)}$.*

This conjecture seems plausible for two reasons: (i) there is a graph on n vertices of pathwidth k having parametric complexity $n^{\Omega(\log k - \log \log n)}$ (Corollary 5); (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 11). These results suggest that we might be very close to resolving Conjecture 46.

We also used alternation-free sequences as a combinatorial way to view parametric shortest paths. Although Martina's counterexample [Kuc18, Example 3.11] shows that there exist infeasible alternation-free sequences, the following question is interesting: how many paths we need to remove from an alternation-free sequence of paths so that the sequence becomes feasible?

Conjecture 47. *If the length of the longest alternation-free sequence of paths in a graph is π, then its parametric complexity is at least $C\pi$, where C is a universal constant.*

If this conjecture is true, then it would imply that alternation-free sequences bound parametric shortest paths from both below and above. However, all known methods employed for assigning linearly varying weights to the edges do not use the fact that the paths form an alternation-free sequence, and it might require considerable insight to resolve Conjecture 47.

Acknowledgments

We would like to 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 Pranabendu Misra for helping us derive the proof for bounded treewidth graphs.

References

[AV18] Nima Anari and Vijay 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.

[BK09] Glencora Borradaile and Philip Klein. An $O(n \log n)$ algorithm for maximum st-flow in a directed planar graph. *J. ACM*, 56(2):9:1–9:30, April 2009.

[Car83a] Patricia J. Carstensen. Complexity of some parametric integer and network programming problems. *Math. Program.*, 26(1):64–75, 1983.

[Car83b] Patricia J. Carstensen. *The complexity of some problems in parametric linear and combinatorial programming*. PhD thesis, University of Michigan, 1983.

[CC16] Chandra Chekuri and Julia Chuzhoy. Polynomial bounds for the grid-minor theorem. *Journal of the ACM (JACM)*, 63(5):40, 2016.

[CFLY10] Sourav Chakraborty, Eldar Fischer, Oded Lachish, and Raphael Yuster. Two-phase algorithms for the parametric shortest path problem. In *27th International Symposium on Theoretical Aspects of Computer Science-STACS 2010*, pages 167–178, 2010.

[CHKM17] José R. Correa, Tobias Harks, Vincent J. C. Kreuzen, and Jannik Matuschke. Fare evasion in transit networks. *Operations Research*, 65(1):165–183, 2017.

[Csa75] Laszlo Csanky. Fast parallel matrix inversion algorithms. In *Foundations of Computer Science, 1975., 16th Annual Symposium on*, pages 11–12. IEEE, 1975.

[Eri10] Jeff Erickson. Maximum flows and parametric shortest paths in planar graphs. In *Proceedings of the twenty-first annual ACM-SIAM symposium on Discrete Algorithms*, pages 794–804. Society for Industrial and Applied Mathematics, 2010.

[Gus80] Daniel Mier Gusfield. *Sensitivity analysis for combinatorial optimization*. PhD thesis, University of California, Berkeley, 1980.

[Gus83] Daniel Mier Gusfield. Parametric combinatorial computing and a problem of program module distribution. *Journal of the ACM (JACM)*, 30(3):551–563, 1983.

[Kas67] Pieter Kasteleyn. Graph theory and crystal physics. *Graph theory and theoretical physics*, pages 43–110, 1967.

[KO81] Richard M. Karp and James B. Orlin. Parametric shortest path algorithms with an application to cyclic staffing. *Discrete Applied Mathematics*, 3(1):37–45, 1981.

[Kuc18] Martina Kuchlbauer. Parametric combinatorial optimization problems and their complexity. Master's thesis, Technische Universitt Mnchen, Munich, Bavaria, 80333, Germany, 2018.

[Mat02] Jiří Matoušek. *Lectures on discrete geometry*, volume 212. Springer New York, 2002.

[MS01] Ketan Mulmuley and Pradyut Shah. A lower bound for the shortest path problem. *J. Comput. Syst. Sci.*, 63(2):253–267, 2001.

[Mul99] Ketan Mulmuley. Lower bounds in a parallel model without bit operations. *SIAM Journal on Computing*, 28(4):1460–1509, 1999.

[Nik09] Evdokia Nikolova. *Strategic algorithms.* PhD thesis, Massachusetts Institute of Technology, 2009.

[NKBM06] Evdokia Nikolova, Jonathan A. Kelner, Matthew Brand, and Michael Mitzenmacher. Stochastic shortest paths via quasi-convex maximization. In *European Symposium on Algorithms*, pages 552–563. Springer, 2006.

[San18] Piotr Sankowski. NC algorithms for weighted planar perfect matching and related problems. In *45th International Colloquium on Automata, Languages, and Programming (ICALP 2018)*, volume 107 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 97:1–97:16, Dagstuhl, Germany, 2018. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik.

[Sha01] Pradyut Shah. *Lower bounds for parallel algorithms.* PhD thesis, The University of Chicago, 2001.

A The PRAM lower bounds

Mulmuley & Shah's [MS01] Theorem 1.4 states the following.

Theorem 48. *The Shortest Path Problem cannot be computed in $o(\log n)$ steps on an unbounded fan-in PRAM without bit operations using $\mathsf{poly}(n)$ 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 [Sha01]) 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 Theorem 6, like Mulmuley & Shah's proof of the corresponding theorem [MS01, Theorem 1.4], is based on the following (see [MS01, Theorem 1.1]).

Theorem 49. *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 49 for *bounded fan-in* PRAMs is established in Mulmuley [Mul99, Theorem 3.3]; Mulmuley & Shah [MS01] 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 [Sha01, Page 36] for an informal justification).

B Thin grids: an application of alternation-free sequences

In this section, we will see that alternation-free sequences can be used to derive upper bounds on the parametric shortest path complexity for certain graph classes. Specifically, we consider a sub-class of planar graphs known as grid graphs.

Definition 50. *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) $((i_1, j_1), (i_2, j_2)) \in E(\Upsilon_{p,q})$[6] *if and only if* $(i_1 = i_2$ *and* $j_2 = j_1 + 1)$ *or* $(j_1 = j_2$ *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. ♢

Let $\varphi^{\mathrm{gr}}(p, q, b)$ 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 b. As a consequence of our main result, $\varphi^{\mathrm{gr}}(n, n, O((\log n)^3)) \geq n^{\Omega(\log n)}$ (Corollary 3). 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 $\varphi^{\mathrm{gr}}(1, n, \infty) \leq 1$ and $\varphi^{\mathrm{gr}}(2, n, \infty) \leq n$ trivially. The problem becomes nontrivial for $3 \times n$ grids. We have the following result.

Theorem 51. $\varphi^{\mathrm{gr}}(3, n, \infty) \leq 5n$.

Proof. Our proof is via on 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 $\{v_1, v_2, \ldots, v_n\}$, 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

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

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 $1 \leq i \leq j \leq n$, 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 $\mathcal{P} = (P(i_1,j_1), P(i_2,j_2), \ldots, P(i_T,j_T))$. 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 $k \in \{1,2,\ldots,T\}$, 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 ($i_r = i_k$ 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))[7], 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 52 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 $|\mathcal{P}| = T \leq 5n$. This completes the proof. □

Claim 52. *Let v_m and $\mathcal{P}$ be as defined in the proof of Theorem 51. Then the incoming and outgoing edges of v_m in $\mathcal{P}$ can each change at most two times.*

Proof. First, we will show that the incoming edge to v_m in $\mathcal{P}$ can change at most two times. Let $\mathsf{pred}_r(v_m)$ be the predecessor of v_m on the path $P(i_r, j_r)$. Since v_m has in-degree 2 (for $m > 1$), $\mathsf{pred}_r(v_m)$ 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 path after that in $\mathcal{P}$. Let us now make this formal.

Suppose there exist four paths $P(i_a, j_a), P(i_b, j_b), P(i_c, j_c), P(i_d, j_d)$ in $\mathcal{P}$ with $a < b < c < d$ such that $\mathsf{pred}_a(v_m) = \mathsf{pred}_c(v_m) = v_{m-1}$ and $\mathsf{pred}_b(v_m) = \mathsf{pred}_d(v_m) = x$. 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.

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

Similarly, it can also be shown that the outgoing edge from v_m in $\mathcal{P}$ can change at most two times. This completes the proof of the claim. $\square$

A simple induction on the grid size shows the following generalization of Theorem 51.

Theorem 53. *For $3 \le p \le q$, we have $\varphi^{\text{gr}}(p, q, \infty) \le O(q(\log q)^{p-3})$.*

Proof. The proof is by induction on p. For the base case ($p = 3$), Theorem 51 automatically implies that $\varphi^{\text{gr}}(3, q, \infty) \le O(q)$. For the inductive case, fix a value of p (where $4 \le p \le q$), and assume that $\varphi^{\text{gr}}(p', q, \infty) \le O(q(\log q)^{p'-3})$ for all $3 \le p' < p$. Now $\Upsilon_{p,q}$ has p rows and q columns; let the vertices in its $\lceil \frac{q}{2} \rceil$-th column be $\{u_1, u_2, \ldots, u_p\}$, 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 subsequences[8] $\mathcal{P}_1, \mathcal{P}_2, \ldots, \mathcal{P}_p$, and provide an upper bound for each. The sum of these p upper bounds is clearly an upper bound on $|\mathcal{P}|$.

Let $\mathcal{P} = \mathcal{P}_1 \uplus \mathcal{P}_2 \uplus \cdots \uplus \mathcal{P}_p$, 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 $\lceil \frac{q}{2} \rceil$-th column that lies on P. For each $\mathcal{P}_i$, we have

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

We are now ready to provide an upper bound for $\varphi^{\text{gr}}(p, q, \infty)$.

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

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

$$\begin{aligned}
\varphi^{\text{gr}}(p, q, \infty) &\le 2\,\varphi^{\text{gr}}(p, r, \infty) + 2 \sum_{i=1}^{p-1} \varphi^{\text{gr}}(i, r, \infty) && (\text{where } r = q - \lfloor q/2 \rfloor + 1) \\
&\le 2\,\varphi^{\text{gr}}(p, r, \infty) + 2c_1 \sum_{i=1}^{p-1} r(\log r)^{i-3} && (\text{applying induction; here } c_1 \text{ is a constant}) \\
&\le 2\,\varphi^{\text{gr}}(p, r, \infty) + 2c_1 r(c_2(\log r)^{p-4}) && (\text{the constant } c_2 \text{ handles lower order terms}) \\
&\le 2\,\varphi^{\text{gr}}(p, r, \infty) + O(q(\log q)^{p-4}).
\end{aligned}$$

[8] A subsequence of an alternation-free sequence is also alternation-free.

Since r is roughly $q/2$, evaluating this final recurrence gives $\varphi^{\mathrm{gr}}(p, q, \infty) \leq O(q(\log q)^{p-3})$. This completes the proof. □

Remark. Theorem 53 only helps for small values of p, that is, when $p \leq \frac{(\log q)^2}{\log \log q}$. For larger values of p, the generalized upper bound of Gusfield [Gus80, Gus83] gives a much better upper bound.

ECCC ISSN 1433-8092
https://eccc.weizmann.ac.il