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

Parametric Shortest Paths in Planar Graphs

Kshitij Gajjar* and Jaikumar Radhakrishnan†

Tata Institute of Fundamental Research, Mumbai

December 3, 2018

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 increases, the 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 an 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 $n^{\log n+O(\alpha(n)^d)}$, where $\alpha(n)$ is the slowly growing inverse Ackermann function. This upper bound arises from 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 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 complexity of the shortest path problem*, denoted by $\varphi(n, b(n))$, be the maximum possible number of break points in $\mathcal{C}(\lambda)$ for a graph with n vertices, where the coefficients in the weights of the edges is bounded by $b(n)$. Let $\varphi^{\mathrm{pl}}(n, b(n))$ be the complexity when the graphs are restricted to be planar. Here is our main result.

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

*kshitij.gajjar@tifr.res.in

†jaikumar@tifr.res.in

ISSN 1433-8092

Before this work similar results were known for general graphs. Carstensen [Car83b] showed that $\varphi(n,\infty) = n^{\Omega(\log n)}$; her result was simplified and extended by Mulmuley & Shah [MS01], who showed that $\varphi(n,(\log n)^3) = n^{\Omega(\log n)}$. Carstensen [Car83a] also presented a matching upper bound argument, $\varphi(n,\infty) = n^{\log n + O(1)}$, which she attributed to Dan Gusfield [Gus80] (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]). $\varphi^{\text{pl}}(n,\infty) = n^{O(1)}$.

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

1.1 Significance of the main result

From their result, $\varphi(n,(\log n)^3) = n^{\Omega(\log n)}$, 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 3. *There exist constants $\alpha > 0$ and $\epsilon > 0$, 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 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 their counterparts 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] 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 graph by considering a parametric shortest paths tree (see Karp &

Orlin [KO81]) in the dual graph, and showing that the tree can undergo only a limited number of changes. Erickson 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.

- 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.

1.2 Previous approaches and our approach

It is worthwhile to examine earlier approaches towards solving Nikolova's conjecture Conjecture 2, and why our approach succeeded where earlier attempts failed.

Previous approaches: We refer to two earlier efforts in resolving this conjecture. 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 linear in n.

Our approach: It is instructive[1] 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 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 4 (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, $p_j[u : v] = q$.*

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

The length of the longest sequence of alternation-free paths is an upper bound on $\varphi^{\mathrm{pl}}(n,\infty)$. Let $f(n,m)$ be the length of the longest sequence of alternation-free paths in the layered graph $G[n,m]$[2]; let $f^{\mathrm{pl}}(n,m)$ be the length of the longest sequence of alternation-free paths in any planar subgraph of $G[n,m]$ (with vertices s and t included). Using the alternation-free property one observes $f(n,1)=n$ and $f(n,2^k-1)\leq 2nf(n,2^{k-1}-1)$, which yields $f(n,2^k-1)\leq \frac{1}{2}(2n)^{k-1}$, implying $\varphi(n,\infty)=n^{O(\log n)}$. The 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 2.1, we present a construction that 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, this construction when adapted to planar graphs yields the following.

Theorem 5. $f^{\mathrm{pl}}(n,(n-1)2^k)\geq n^k$.

In Subsection 2.3, we present this construction 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. For example, Kuchlbauer [Kuc18, Example 3.11] presents a planar (grid) 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 paths as shortest paths.

Our refutation of Nikolova's conjecture is based on the construction of Mulmuley & Shah [MS01]. The Mulmuley-Shah construction uses an intricate inductive argument involving the composition of dense bipartite graphs. These bipartite graphs contain large complete bipartite graphs, and are therefore highly non-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 with up to n^2 edges and the original weight is carefully distributed between them. 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 don't 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 4.1 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.

2 Alternation-free paths in graphs

In this section, we outline the construction of a planar graph with a large number of alternation-free paths. Note that this graph is similar to (and inspired by) earlier examples of graphs

[2] $G[n,m]$ has m layers with n vertices in each layer.

with alternation-free paths [Car83b, MS01]. For this, 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 the standard lexicographic order, correspond to a sequence of alternation-free paths.

2.1 Alternation-free paths and alternation-free sequences

We first present a version of alternation-free paths for non-planar graphs, and then refine it to obtain another for a related planar graph. Consider a 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 $\mathbb{Z}_n = \{0, 1, \ldots, n-1\}$. In the non-planar version, we add all edges from a layer to the next. We refer to the resulting graph as $G^{\mathrm{npl}}[n, m]$:

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

Thus $L_{j-1} \cup L_j$ is a complete bipartite graph for $j = 1, 2, \ldots, m-1$ in $G^{\mathrm{npl}}[n, m]$. 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 $G^{\mathrm{pl}}[n, m]$:

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

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-1} goes around the surface of the cylinder to the 0-th vertex in layer L_j). 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 6 (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 an alternation-free sequence of words, 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 7 (Alternation-free sequence of words)**.** *Let S be sequence of ℓ words from $\mathbb{Z}_n^m$, that is, $S \subseteq (\mathbb{Z}_n^m)^\ell$. 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 8 (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 special case $v = m$ in the second condition of Definition 7 is used to verify that there is no alternation between pairs of the form (u, m). □

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 size n^ℓ) over $\mathbb{Z}_n$ and $\{0, 1\}$ respectively. Let us 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). Considering the same example as in the previous paragraph, $\hat{X}[114] = (110\ 100\ 111\ 100\ 000\ 100\ 111\ 100)$. Now we will show that X and $\hat{X}$ are alternation-free.

Lemma 9. *Suppose $S \in (\mathbb{Z}_n^m)^\ell$ is an alternation-free sequence of ℓ words in $\mathbb{Z}_n^m$. Then,*

(*a*) *$S \downarrow j$ 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 $\mathbb{Z}_n^{2m}$.*

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 (s, t) $(0 \leq s < t \leq 2m)$, 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) $(0 \leq a < b < c \leq n\ell)$ between (s, t) $(0 \leq s < t \leq 2m)$. 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 $t - s \geq 2$ (as observed above). 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 10. *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 9 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 9, 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.

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 11. *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) $\hat{S} \downarrow \hat{\jmath}$ 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. 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) $(0 \leq a < b < c \leq \ell - 1)$ between (s, t) $(0 \leq s < t \leq 2m(n-1))$. 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 $\sigma_a[s] = \sigma_c[s] \neq \sigma_b[s]$, and similarly that $\sigma_a[t-1] = \sigma_c[t-1] \neq \sigma_b[t-1]$. 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) $(0 \leq a < b < c \leq n\ell - 1)$ between (s, t) $(0 \leq s < t \leq 2m(n-1))$; 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 $\sigma_a[s:t] \neq \sigma_c[s:t]$, and there can be no alternation. So we may assume that there is no full odd numbered in the range $\{s, s+1, \dots, t-1\}$. To complete the proof, we now use two crucial facts about our encoding. First, in the sequence of words $(\hat{0}, \hat{1}, \dots, \widehat{n-1})$, 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 Upper bound for polynomial edge weights

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 $\varphi_d(n, b(n))$ 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 with coefficients bounded by $b(n)$.

Theorem 12. *For all fixed d, we have $\varphi_d(n, \infty) = n^{\log n + O(\alpha(n)^d)}$, where $\alpha(n)$ is the extremely slow growing inverse Ackermann function.*

Proof. We adapt to our setting an argument due to Dean (see Nikolova [Nik09, Page 86]). Let $f(n, m)$ be the maximum length of a sequence of shortest paths, 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 $2nf(n, 2^{k-1})$ distinct paths in the sequence σ. 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 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 $N \le 2nf(n, 2^{k-1})$. Bounds known for Davenport-Schinzel sequences (see Matoušek [Mat02, Page 173]) imply that the maximum length of a sequence of shortest paths is at most $N 2^{\alpha(N)^d}$ (for all large N). Since $N \le n^n$ (a coarse upper bound on the total number of paths in any n-vertex graph), we have $2^{\alpha(N)^d} \ll 2^{(2\alpha(n))^d}$. Thus, $f(n, 2^k) \le N \cdot 2^{\alpha(N)^d} \le 2nf(n, 2^{k-1}) \cdot 2^{(2\alpha(n))^d}$, which yields $f(n, 2^k) \le (2n)^k 2^{k(2\alpha(n))^d}$. Our theorem follows from this by taking $k = \lceil \log n \rceil$. □

4 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 [MS01], which in turn was based on the construction of Carstensen [Car83b]. These earlier constructions proceed by induction, wherein each level of induction increases the number of vertices by a constant factor and the number of breakpoints in the lower envelope by a factor n. After m steps of induction, we obtain a graph $G_{m,n}$ with $\mathsf{poly}(n) \cdot \exp(m)$ vertices and n^m paths. Figure 2 depicts how the graph $G_{m,n}$ is put together from two copies of $G_{m-1,n}$ and one copy of the form $G_{m-1,2n-1}$. The edge weights in the constituent graphs are carefully chosen, but are not important to our top-level view. The only new component added in this level of recursion is the part labelled LINK.

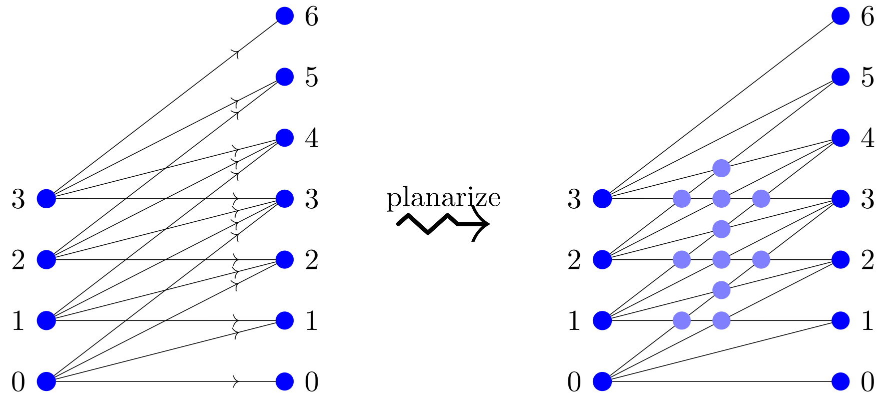

Figure 1: The LINK *gadget* $\mathbf{L}(B,n)$ *for* $B=4, n=3$ *and its planarization*

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 5, we show that the new edge weights in LINK also result in a large number of break points. In Subsection 4.1, 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.

4.1 The linking gadget

A linking gadget $\mathbf{L}(B,n)$ is a bipartite graph $G(U,V,E,(w_e : e \in E))$ with $U = \{0,1,\dots,B-1\}$, $V = \{0,1,\dots,B+n-1\}$, $E = \{(b,b+r) : b \in U, r = 0,1,\dots,n\}$. 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 $\mathbb{R}^2$ 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 1). 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 most n. Thus the horizontal distance traversed by an edge $((x_1,y_1),(x_2,y_2))$ of G^{pl} (that is, x_2-x_1), can be written as a *non-zero fraction with denominator at most* n^2. We use this observation in Claim 15 below.

Definition 13. *We say that* G^{pl} faithfully simulates G *if for all* $(b,j)\in U\times V$*:*

(*i*) *if* $b \leq j \leq b+n$*, 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* $b \leq j \leq b+n$*, 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* $w_{b,j} + 1$,

(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 paths in G^{pl} *should behave like edges in* G. ◊

Lemma 14. *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 $b \leq j \leq b + n$. 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 15. $\mathcal{C}(Q) - \mathcal{C}(P) \geq n^{-2}K - 2 \max_{e \in E(G)} |J(e)|$.

Proof of claim. Let Q consist of vertices $q_0 = (x_0, y_0), q_1 = (x_1, y_1), q_2 = (x_2, y_2), \ldots, q_t = (x_t, y_t)$, where $(x_0, y_0) = (0, b)$ and $(x_t, y_t) = (1, j)$. 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 $\beta_i = x_i - x_{i-1}$. Then for $i \in 1, 2, \ldots, t$, we have

$$\beta_i \geq n^{-2}; \quad \text{(by the observation above)}$$

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

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

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

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

$$\mathcal{C}(Q) - \mathcal{C}(P) \geq -2|\max_e J(e)| + K\left(\mathbb{E}\left[\frac{r_{\mathbf{i}}^2}{2}\right] - \frac{r^2}{2}\right) + \left(\frac{K + 2L\lambda}{2}\right)(\mathbb{E}\left[r_{\mathbf{i}}\right] - r) \tag{16}$$

$$\geq -2|\max_e J(e)| + K\left(\mathbb{E}\left[\frac{r_{\mathbf{i}}^2}{2}\right] - \frac{r^2}{2}\right) \tag{17}$$

$$\geq -2|\max_e J(e)| + \frac{K}{2}\mathbf{var}[r_{\mathbf{i}}]. \tag{18}$$

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,

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

Combining this with (18) proves Claim 15. □

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. □

5 Proof of the main result

We now prove Theorem 1. The proof is by induction. Intuitively, we start off 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. (Throughout our proof, n is a fixed number.)

5.1 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. In this section, we place the framework by describing the intervals inductively. The intervals depend on a parameter N, defined as

$$N = n^2. \tag{19}$$

Then, for $m = 0, 1, \ldots, \log n$ and $j = 0, 1, \ldots, n^m - 1$, we define $\alpha(j, m) \in \mathbb{R}$ inductively; these points will be used to define intervals.

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

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

Intervals: For $m \geq 0$ and $0 \leq j < n^m$, let $\mathcal{I}(j, m) = [\alpha(j, m) + 1, \alpha(j, m) + N - 1]$. Since $0 \leq r < n \ll N$, we have that for $m \geq 1$, the interval $\mathcal{I}(j, m)$ is a subset of the interval $N \cdot \mathcal{I}(d, m-1)$, which is obtained by stretching $\mathcal{I}(d, m-1)$ by a factor N.

5.2 Inductive construction of graphs

Our induction depends on several parameters which impose constraints on the layered, weighted, planar graphs we construct. The parameter B denotes the number of vertices in the first (input) layer of this graph, 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 $B = 1$, call the unique input vertex s, and connect all the vertices in the last layer to a new vertex t using edges with weight 0, so that we have pristine s-t paths as promised.) The number $D \in \mathbb{R}$ is used to determine the weights of the edges. Finally, the induction parameter m (this is the same induction parameter which is used to define the intervals) helps ensure that the number of break points in the cost of the shortest path is large. See Figure 3 for a nice step-by-step visualization of this construction.

The predicate Φ: 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 at most $(3^{m+1} - 1)(B + mn)^4$ vertices, B input vertices, and paths P_{bj} (for $b = 0, 1, \ldots, B-1$ and $j = 0, 1, \ldots, n^m - 1$) such that

(*i*) 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;

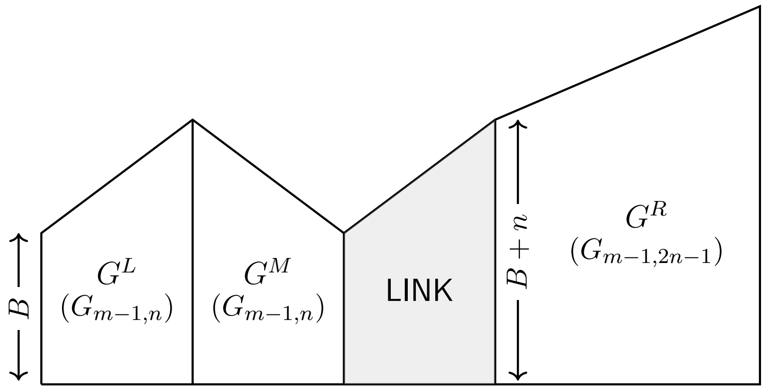

Figure 2: The graph $G_{m,n}$ is obtained by composing G^L, G^M, the linking gadget, and G^R

(*ii*) for all b and j, we have $\mathcal{C}(P_{bj})(\lambda) = \mathcal{C}(P_{0j})(\lambda) + bD\alpha(j,m)$;

(*iii*) for all j, the paths in the list $(P_{bj} : b = 0, 1, \ldots, B-1)$ are vertex-disjoint;

(*iv*) 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 4.1 and Lemma 14 to keep our graphs planar.

Lemma 20 (Main lemma)**.** *For all B, D and $m \geq 0$, the predicate $\Phi(B, D, m)$ holds.*

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

Proof of Theorem 1. By Lemma 20, taking $B = 1$, $D = 0$ and $m = \lfloor \log n \rfloor$, we conclude that the predicate $\Phi(1, 0, \lfloor \log n \rfloor)$ holds. The number of vertices in the corresponding graph $G = (V, E)$ is at most

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

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; so the cost of the shortest s-t path in this graph has $n^{\lfloor \log n \rfloor}$ break points. We also show that the bit lengths of coefficients involved in this construction are at most $C(\log n)^3$ for some constant C. Let ν be a large positive integer. Let n be the largest integer such that $6n^8 + 1 \leq \nu$ and $C(\log n)^3 \leq (\log \nu)^3$. Note $n = \nu^{\Omega(1)}$. Using the construction above (adding dummy isolated vertices if necessary), we obtain a graph on ν vertices, whose edge weights have coefficients bounded by $(\log \nu)^3$, and in which the cost of the shortest path has $\nu^{\Omega(\log \nu)}$ break points. $\square$

The remainder of this paper is devoted to proving our main lemma.

Proof of Lemma 20. We will use induction on m. For the base case ($m = 0$), let G consist of B disjoint edges, each with weight 0, leaving the B input vertices. The only choice for j in this case is $j = 0$ (since j varies from 0 to $n^m - 1$), and the requirements for Φ are easily verified. (To verify (ii), recall that $\alpha(0,0) = 0$.)

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

$$K_L = 100(|D|^2 + 1)N^{m+4}B^2 \tag{21}$$

$$K_R = 10(|D| + 1)N^3 B \tag{22}$$

$$D_L = \frac{N}{2K_L}\left(D - \frac{K_R}{N}\right) \tag{23}$$

$$D_R = 1 \tag{24}$$

The assignments to these constants may seem mysterious but they will be justified by the claims that follow. Let G^L be the graph corresponding to the induction hypothesis $\Phi(B, D_L, m-1)$; we denote the corresponding $B \times n^{m-1}$ 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. Thus G^M has B vertices in its last layer (see Figure 2); let $(P_{bj})^{\text{rev}}$ be the *reverse* of the path P_{bj} (thus, P_{bj} starts at the vertex b of the first layer of G^M and ends at the vertex b of the last layer of G^M). Let G^R be the graph corresponding to the induction hypothesis $\Phi(B+n, D_R, m-1)$; we denote the corresponding $B \times n^{m-1}$ paths by P^R_{bj} where $0 \leq b < B$ and $0 \leq j < n^{m-1}$.

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. We replace the weight function $w_e(\lambda)$ by $K_L \cdot w_e(\lambda/N)$ for each edge e in G^L and G^M, and replace the weight function $w_e(\lambda)$ by $K_R \cdot w_e(\lambda/N)$ for each edge e in G^R. In essence, we are *scaling* (by factors K_L and K_R) and *stretching* (by a factor N) our already existing solutions for G^L, G^M and G^R so that together they can form a solution for G. Let $\mathbf{L}(B, n)$ be the non-planar linking gadget with edge weights

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

and let $\mathbf{L}^{\text{pl}}(B, n)$ be its planarized version. The graph G obtained by composing G^L, G^M, $\mathbf{L}^{\text{pl}}$ and G^R is shown in Figure 2. Since G^L, G^M and G^R are planar by induction, and $\mathbf{L}^{\text{pl}}(B, n)$ is planar, G is also planar.

Before we proceed further, let us verify that for our choice of parameters, $\mathbf{L}^{\text{pl}}$ faithfully simulates its non-planar counterpart. Invoke Lemma 14 with $J(b, b+r) = NDrb$, $K = K_R$ and $L = -K_R/N$. For the setting of K_R in (22), we have

$$n^2 \left(1 + 2 \max_{e \in E(G)} |J(e)|\right) \leq n^2\ (1 + 2N|D|nB) \leq 4(|D|+1)N^{2.5}B \leq K_R,$$

so the condition $K \geq 2n^2\,(1 + \max_{e \in E(G)} |J(e)|)$ of Lemma 14 holds.

To show that $\Phi(B, D, m)$ holds, we need to exhibit $B \times n^m$ paths in G and verify that conditions (i)–(iv) hold. For $0 \leq j < n^m$, write $j = nd + r$ with $0 \leq d < n^{m-1}$ and $0 \leq r < n$; then for $0 \leq b < B$, let

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

where $\texttt{link}(b, b+r+1)$ is the unique shortest path (the straight line) in $\mathbf{L}^{\text{pl}}$ connecting vertex b in the last layer of G^M to vertex $b+r+1$ in the first layer of G^R. Since the planarization of

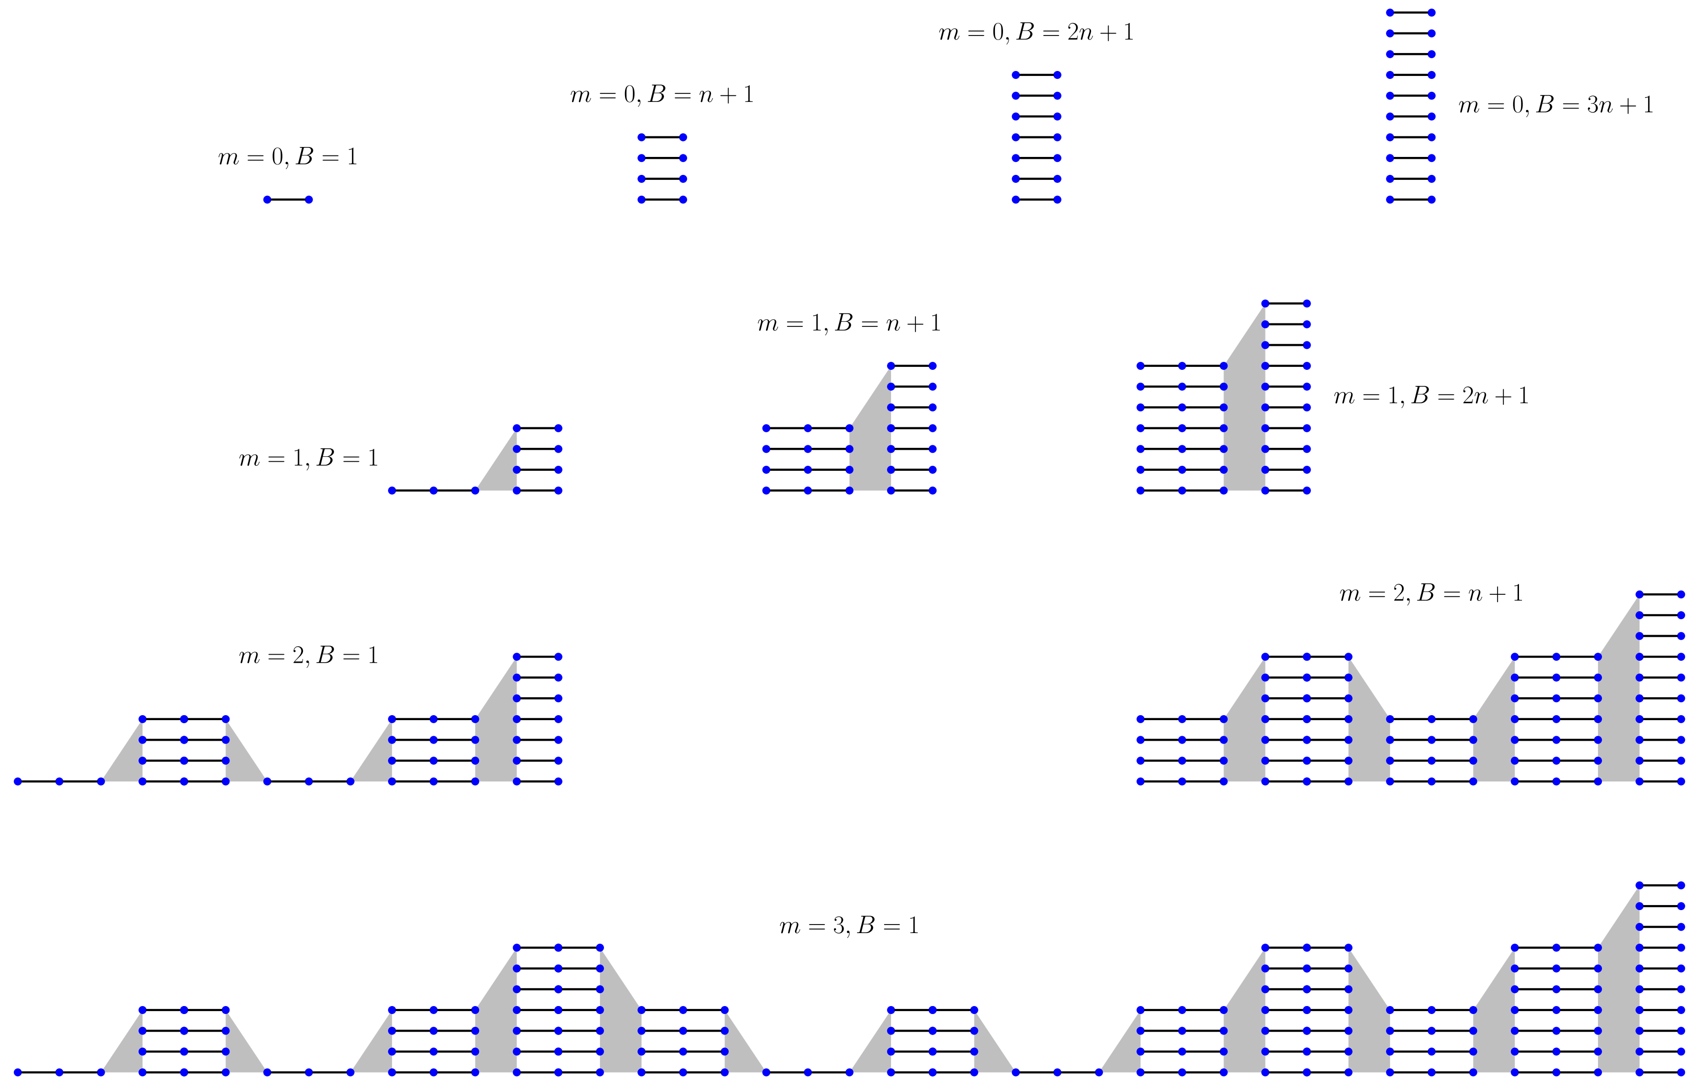

Figure 3: An instantiation of the graph $G_{m,n}$ for $m = 3, n = 3$

the linking gadget adds at most $(B+n)^4$ new vertices, the number of vertices in the planarized version of G is at most

$$2(3^m-1)(B+(m-1)n)^4+(B+n)^4+(3^m-1)(B+n+(m-1)n)^4 \le (3^{m+1}-1)(B+mn)^4.$$

The requirements (iii) and (iv) are straightforward to verify. We now verify (ii).[3]

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

Substitute $b=0$ to get

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

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

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

as required. To complete the verification of (i), 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.

In Claim 25, we track paths from an input vertex as they travel through G^L and G^M. In Claim 26, we analyze how such paths continue through $\mathbf{L}^{\text{pl}}$ and G^R. Fix j $(0 \le j < n^m - 1$, say $j = nd + r$, for $0 < d < n^{m-1}$ and $0 \le r < n)$ and a $\lambda \in \mathcal{I}(j,m)$. Note that since $\lambda \in \mathcal{I}(j,m)$, we have $\lambda/N \in \mathcal{I}(d,m-1) = [\alpha(d,m-1)+1, \alpha(d,m-1)+N-1]$.[4]

Claim 25. *Let Q be a path from the input vertex b to the last layer of $G^L \circ G^M$ (note that $P^L_{bd} \circ (P^L_{bd})^{\text{rev}}$ is one such path). If $Q \neq P^L_{bd} \circ (P^L_{bd})^{\text{rev}}$, then*

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

[3] Since 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.

[4] In fact, our recursive definition of $\alpha(j,m)$ has $(r+1)N$ instead of rN 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.

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,

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

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 $Q \neq P^L_{bd} \circ (P^L_{bd})^{\text{rev}}$, 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 part (i) 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 (21) we have $K_L B D_L \alpha(d, m-1) \ll K_L/10$. This completes the proof of Claim 25. □

Since K_L is positive, Claim 25 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 $G^L \circ G^M$. We need to argue that the overall shortest path must be an extension of this.

Claim 26. *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 $\mathbf{L}$ 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*

$$\mathcal{C}(P)(\lambda) - \mathcal{C}(\texttt{link}(b, b+r+1) \circ P^R_{b+r+1,d})(\lambda) \geq 1\,. \tag{27}$$

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

$$Z_0 \gg Z_1 \gg \cdots \gg Z_r \gg Z_{r+1} \ll Z_{r+2} \ll \cdots \ll Z_n, \tag{28}$$

where we use $\gg$ and $\ll$ to suggest that there is a large gaps between the quantities. 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

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

Thus,

$$\begin{aligned}
Z_k - Z_{k-1} &= NDb + \frac{K_R}{N}(kN-\lambda) + K_R D_R\alpha(d,m-1) && (29)\\
&= NDb + \frac{K_R}{N}(kN + N\alpha(d,m-1) - \lambda) \qquad \text{(recall } D_R = 1) && (30)\\
&= NDb + \frac{K_R}{N}(\alpha(k-1,m) - \lambda). && (31)
\end{aligned}$$

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

$$\begin{aligned}
\alpha(k-1,m) - \lambda &\in [\alpha(k-1,m) - \alpha(r,m) - N + 1, \alpha(k-1,m) - \alpha(r,m) - 1] && (32)\\
&= [(k-(r+1))N - N + 1, (k-(r+1))N - 1]. && (33)
\end{aligned}$$

Thus, for $k = 1, 2, \ldots, r+1$, we have $\alpha(k-1,m) - \lambda \leq -1$ and for $k = r+2, \ldots, n$, we have $\alpha(k-1,m) - \lambda \geq +1$. Returning to (31) with this, we obtain

$$\begin{aligned}
Z_k - Z_{k-1} &\leq NDb - \frac{K_R}{N} \qquad \text{for } k = 1, \ldots, r+1, \text{ and} && (34)\\
Z_k - Z_{k-1} &\geq NDb + \frac{K_R}{N} \qquad \text{for } k = r+2, \ldots, n. && (35)
\end{aligned}$$

Since $K_R \gg N^2Db$, the RHS of (34) is negative and the RHS of (35) is positive. This confirms (28) and establishes Claim 26. □

We are now in a position to establish (i) and complete the induction. By Claim 25, if the shortest path from b does not follow $P^L_{b,d} \circ (P^R_{b,d})^{\mathrm{rev}}$, then the increase in cost is at least $K_L/2$. The shortest path from an input vertex of $\mathbf{L}^{\mathrm{pl}}$ to the last layer of G^R has cost at most

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

So any compensation from $\mathbf{L}^{\mathrm{pl}} \circ G^R$ is at most $K_L/10$. Thus the shortest path in G must follow the prescribed route in $G^L \circ G^M$ until it arrives at the first layer of $\mathbf{L}^{\mathrm{pl}}$ (or it already incurs an increase in cost of $K_L/2 - K_L/10 \gg 1$, regardless of what route it takes in $\mathbf{L}^{\mathrm{pl}} \circ G^R$). Claim 26 now confirms that it must continue by taking the edge $\mathtt{link}(b, b+r+1)$ and $P^R_{b+r+1,d}$; any deviation from this path will incur an increase in cost of at least 1.

Space Complexity: In the description above, we did not explicitly keep track of the growth of the coefficients involved in the edge weights. Since K_L is the largest of the four constants (K_L, K_R, D_L, D_R) involved in computing the edge weights, it is sufficient to track the growth of K_L. In terms of n, we have $N = n^2$ and $B \leq n \log n$ (all of these are $\mathsf{poly}(n)$; we also need not consider D separately since it is initialized to 0).

Note that in the m-th level of recursion, K_L grows by a factor of at most n^{mc}, where c is some constant (21). After m steps, K_L has grown as large as $n^c \cdot n^{2c} \cdot n^{3c} \cdots n^{mc} = n^{O(m^2)}$. Since our induction terminates at $m = \log n$, we have $K_L \leq n^{O((\log n)^2)} = 2^{O((\log n)^3)}$. Thus each edge weight can be stored using $O((\log n)^3)$ bits of memory, completing the proof of Lemma 20. $\square$

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, and Suhail Sherif for helping us simulate the problem using a computer program.

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 June Carstensen. *The complexity of some problems in parametric linear and combinatorial programming*. PhD thesis, University of Michigan, 1983.

[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.

[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 Velinova 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.

[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 36. *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 $O(\log^3 n)$.*

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 $O((\log n)^3)$ 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 3, like Mulmuley & Shah's proof of the corresponding theorem [MS01, Theorem 1.4], is based on the following (see [MS01, Theorem 1.1]).

Theorem 37. *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 37 for *bounded fan-in* PRAMs is established in Mulmuley [Mul99, Theorem 3.3]; Mulmuley and 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).

ECCC ISSN 1433-8092

https://eccc.weizmann.ac.il