

Improved Compression of the Okamura-Seymour Metric

Shay Mozes¹, Nathan Wallheimer², and Oren Weimann²

¹ The Interdisciplinary Center Herzliya, Israel, smozes@idc.ac.il

² University of Haifa, Israel, nathanwallh@gmail.com, oren@cs.haifa.ac.il

Abstract

Let $G = (V, E)$ be an undirected unweighted planar graph. Consider a vector storing the distances from an arbitrary vertex v to all vertices $S = \{s_1, s_2, \dots, s_k\}$ of a single face in their cyclic order. The *pattern* of v is obtained by taking the difference between every pair of consecutive values of this vector. In STOC'19, Li and Parter used a VC-dimension argument to show that in planar graphs, the number of *distinct* patterns, denoted $p_{\#}$, is only $O(k^3)$. This resulted in a simple compression scheme requiring $\tilde{O}(\min\{k^4 + |T|, k \cdot |T|\})$ space to encode the distances between S and a subset of terminal vertices $T \subseteq V$. This is known as the Okamura-Seymour metric compression problem.

We give an alternative proof of the $p_{\#} = O(k^3)$ bound that exploits planarity beyond the VC-dimension argument. Namely, our proof relies on cut-cycle duality, as well as on the fact that distances among vertices of S are bounded by k . Our method implies the following:

- (1) An $\tilde{O}(p_{\#} + k + |T|)$ space compression of the Okamura-Seymour metric, thus improving the compression of Li and Parter to $\tilde{O}(\min\{k^3 + |T|, k \cdot |T|\})$.
- (2) An optimal $\tilde{O}(k + |T|)$ space compression of the Okamura-Seymour metric, in the case where the vertices of T induce a connected component in G .
- (3) A tight bound of $p_{\#} = \Theta(k^2)$ for the family of Halin graphs, whereas the VC-dimension argument is limited to showing $p_{\#} = O(k^3)$.

1 Introduction

Planar metric compression. The shortest path metric of planar graphs is one of the most popular and well-studied metrics in computer science. The planar graph *metric compression* problem is to compactly encode the distances between a subset of k terminal vertices so that we can retrieve the distance between any pair of terminals from the encoding. On an n -vertex planar graph $G = (V, E)$, a naïve encoding uses $\tilde{O}(\min\{k^2, n\})$ bits (by either storing the $k \times k$ distance matrix or alternatively by storing the entire graph¹). It turns out that this naïve bound is actually optimal (up to logarithmic factors) for *weighted* planar graphs, as shown by Gavaille et al. [16]. It is important to note that their lower bound applies even when all terminals lie on a single face. The complexity of *unweighted undirected* planar graphs is also well-understood. Gavaille et al. [16] (see also [1]) gave a lower bound of $\Omega(\min\{k^2, \sqrt{k \cdot n}\})$, and Abboud et al. [1] gave a matching upper bound.

¹Naïvely, this takes $O(n \log n)$ bits, but can be done with $O(n)$ bits [4, 10, 26, 30].

If we are willing to settle for *approximate* distances, then there are ingenious compressions requiring only $\tilde{O}(k)$ bits [21, 22, 29]. The problem has also been extensively studied (both in the exact [6, 7, 9, 17, 23, 24] and the approximate [2, 5, 8, 13, 14, 18, 19] settings) for the case where we require that the compression is itself a graph (that contains the terminals and preserves their distances).

The Okamura-Seymour metric compression. An important special case for which tight bounds are not yet known, is when the planar graph is unweighted and undirected and we want to encode the $S \times T$ distances between a set of k source terminals $S = \{s_1, s_2, \dots, s_k\}$ lying consecutively on a single face and a subset of target terminal vertices $T \subseteq V$. A query (v, s_i) to the encoding (with $v \in T$ and $s_i \in S$) returns the v -to- s_i distance.

- When $T = S$, it is possible to exploit the *Unit-Monge* property to obtain an $O(k \log k)$ space encoding with $O(\log k)$ query time [1]. In fact, even if $T \neq S$ the Unit-Monge property implies an (optimal) $\tilde{O}(|T| + k)$ space encoding, as long as the vertices of T lie (not necessarily consecutively) on single face.
- When $T = V$, the MSSP data structure of Eisenstat and Klein [12] gives an $O(n)$ space encoding with $O(\log n)$ query time.
- For arbitrary S, T , Li and Parter [25] recently presented a compression of size $\tilde{O}(\min\{k^4 + |T|, k \cdot |T|\})$ and query time $O(1)$.² This compression is useful algorithmically. In the distributed setting, Li and Parter used it to compute the diameter of a planar graph in $\tilde{O}(\text{poly}(D))$ rounds where D is the graph's diameter. It was also used to develop an exact distance oracle with subquadratic space and constant query time [15].

The Li-Parter compression. At the heart of the Li-Parter compression [25], is the notion of a *pattern*. Let $d(\cdot, \cdot)$ denote the shortest path metric of G . The pattern of a vertex $v \in V$ is the vector

$$p_v = \langle d(v, s_2) - d(v, s_1), d(v, s_3) - d(v, s_2), \dots, d(v, s_k) - d(v, s_{k-1}) \rangle.$$

Since the graph is unweighted, every entry of p_v is in $\{-1, 0, 1\}$ by the triangle inequality. This already gives an efficient way to encode v 's distances to S : Instead of explicitly storing these distances (using $O(k \log n)$ bits), store p_v and $d(v, s_1)$ (using $O(k + \log n)$ bits). This way, any distance $d(v, s_i)$ can be retrieved by

$$d(v, s_i) = d(v, s_1) + \sum_{j=1}^{i-1} p_v[j].$$

The main contribution of Li and Parter in this context is in showing that, while there are overall n patterns in the graph, there are only $O(k^3)$ *distinct* patterns:

Theorem 1 ([25]). *The number of distinct patterns over all vertices of the graph is $O(k^3)$.*

²The actual bound stated in [25] is $\tilde{O}(k^3 \cdot D + |T|)$ where D is the diameter of the graph. The reason for the additional D factor is that they store all possible distance tuples $d_v = \{d(v, s_i)\}_{i=1}^k$ instead of all possible patterns p_v . The reason for the missing k factor is simply a mistake in their paper.

The compression follows easily from the above theorem: Store one table that contains all the distinct patterns of vertices in T , and another table that contains for every $v \in T$ the value $d(v, s_1)$ and a pointer to p_v in the first table. Since there cannot be more than $|T|$ distinct patterns, the size of the first table is $O(\min\{k^4, k \cdot |T|\})$. The size of the second table is $\tilde{O}(|T|)$. The query time is $O(k)$ but can be improved to $O(1)$ by storing precomputed prefix-sums of every pattern (increasing the size of the first table by a logarithmic factor).

The original proof of Theorem 1 [25]. Let us assume that the distinguished face is the infinite face. For convenience, we transform³ the problem so that patterns are binary rather than ternary (i.e. over $\{-1, 1\}$ instead of $\{-1, 0, 1\}$). To this end, we subdivide every edge of the graph to get a new (unweighted) graph G' . In particular, we replace each edge $\{s_i, s_{i+1}\}$ of the infinite face with a dummy vertex w_i and edges $\{s_i, w_i\}, \{w_i, s_{i+1}\}$. For every vertex u of G' , let \hat{p}_u be the pattern of u w.r.t. the set of vertices $S' = \{s_1, w_1, s_2, w_2, \dots, s_k, w_k\}$. Observe that the parity of u -to- s_i distances is different from the parity of u -to- w_j distances, for all i, j 's. Hence, \hat{p}_u is a binary vector (i.e. over $\{-1, 1\}$). Additionally, for every vertex v of G we can retrieve its pattern p_v from \hat{p}_v since $p_v[i] = (\hat{p}_v[2i-1] + \hat{p}_v[2i])/2$. See Figure 1. Hence, we henceforth assume that patterns p_v are over $\{-1, 1\}$ (i.e. we replace p_v with \hat{p}_v). For brevity, we also assume that patterns are of length $k-1$ (rather than $2k-1$).




Figure 1: Before (left) and after (right) the transformation that makes all patterns binary. Every edge is subdivided and a new vertex (in color) is put in the middle. In this example, $p_v = \langle 0, 1, -1, 0, 1, 1, -1 \rangle$ and $\hat{p}_v = \langle 1, -1, 1, 1, -1, -1, 1, -1, 1, 1, 1, 1, -1, -1, -1 \rangle$.

Li and Parter's VC-dimension argument is based on the simple observation that, by planarity, there cannot be two vertices v and u and 4 indices $a < b < c < d$ such that $p_u[a] = -1, p_u[b] = 1, p_u[c] = -1, p_u[d] = 1$ but $p_v[a] = 1, p_v[b] = -1, p_v[c] = 1, p_v[d] = -1$. The reason is that such $(-1, 1, -1, 1), (1, -1, 1, -1)$ patterns correspond to an illegal configuration of shortest paths in planar graphs.

Consider arranging all the patterns as the rows of a binary matrix P . The VC-dimension d of P , is the maximum number of columns in a submatrix of P that contains all possible 2^d rows. The above forbidden configuration implies that there is no submatrix with 4 columns or more that contains all possible rows, hence the VC-dimension of P is at most 3. By the well known Sauer's

³This transformation was suggested by Li and Parter in their STOC'19 talk.

Lemma [27], this means that there are $O((k-1)^3) = O(k^3)$ distinct rows. This is the entire proof.

Limitations of the original proof. It remains an open problem whether the number of distinct patterns in planar graphs is $\Theta(k^3)$ or less (there is a simple $\Omega(k^2)$ lower bound). We do know however that there is no hope of improving $\Theta(k^3)$ using the VC-dimension argument: Consider the following set of sequences over $\{-1, 1\}^{k-1}$:

$$\left\{ (-1)^{x_1} \circ 1^{x_2} \circ (-1)^{x_3} \circ 1^{k-1-x_1-x_2-x_3} \mid x_1 + x_2 + x_3 < k \right\}$$

There is no pair of sequences in this set that contains the forbidden $(1, -1, 1, -1), (-1, 1, -1, 1)$ configuration, and yet its cardinality is $\Theta(k^3)$. This means that any improvement to the $O(k^3)$ bound on the number of distinct patterns in planar graphs would have to further exploit structural properties of planar graphs. In fact, even in the restricted family of Halin graphs, where we know that there are only $\Theta(k^2)$ distinct patterns (see Section 4), the VC-dimension argument is limited to proving $O(k^3)$.

Our results and technique. We develop a new technique for analyzing and encoding the structure of patterns in a planar graph using *bisectors*. The bisector β_i associated with vertex s_i is a simple cycle in the dual graph such that all (primal) vertices on the same side of β_i have the same i 'th bit in their patterns. We show that any two bisectors are arc-disjoint. This implies the following lemma:

Lemma 2. *The patterns of every two adjacent vertices in G differ by at most two bits.*

We then show how to use this property to obtain the following compression (recall that $p_{\#}$ denotes the number of distinct patterns in G):

Theorem 3. *There is an $\tilde{O}(p_{\#} + k + |T|)$ space compression of the Okamura-Seymour metric with $\tilde{O}(n)$ construction time and $\tilde{O}(1)$ query time. Moreover, for the special case where the vertices of T induce a connected component in G , the space is $\tilde{O}(k + |T|)$.*

By plugging $p_{\#} = O(k^3)$ from Theorem 1 (and the trivial compression that stores all $T \times S$ distances) we get an $\tilde{O}(\min\{k^3 + |T|, k \cdot |T|\})$ compression (i.e. a factor k improvement over Li and Parter [25]). Moreover, for the special case where the vertices of T induce a connected component in G , we obtain an optimal $\tilde{O}(|T| + k)$ space encoding. Recall that, prior to our work, this bound was only known (using the Unit-Monge property) when the vertices of T all lie (not necessarily consecutively) on a single face [1]. In fact, even in such setting, our method gives $\tilde{O}(|T| + k)$. Thus, our method strictly dominates the one based on Unit-Monge.

An additional benefit of working with bisectors is that they can be used to bound the number $p_{\#}$ of distinct patterns. We show that every two bisectors can cross only $O(k)$ times. Our proof relies not only on the planar structure, but also on the fact that the distance between any two vertices of S is bounded by k (this property is not used in the VC-dimension argument). The set of all bisectors partitions the plane into regions. All (primal) vertices in the same region have the same pattern because they all lie on the same side of every bisector. Since there are $O(k^2)$ pairs of bisectors, and each pair crosses $O(k)$ times, there are only $O(k^3)$ regions (and hence only $O(k^3)$ distinct patterns). This provides an alternative proof of Theorem 1. We believe that our new technique may prove useful in settling the question of the number of distinct patterns in a planar

graph. In particular, it may be that a similar argument that uses stronger structural properties will be able to show that the partition induces only $O(k^2)$ regions. We demonstrate this potential of our technique in Section 4, where we show such a bound for a family of graphs that includes Halin graphs:

Theorem 4. *The number of distinct patterns over all vertices of a Halin graph is $O(k^2)$. This bound is tight.*

In contrast, the VC-dimension argument is limited to proving $O(k^3)$, even on Halin graphs.

2 Preliminaries

Let $G = (V, E)$ be an unweighted, undirected planar embedded graph. We prefer to think of G as a directed planar graph with a set of arcs \mathcal{A} , such that there is a pair of arcs uv, vu (embedded on the same curve) for every edge $\{u, v\} \in E$. We refer to u and v as the *tail* and *head* of uv , respectively. We refer to uv as the *reverse* of vu , or simply $rev(vu)$. However, we use the term edge whenever the orientation is not important or when we refer to any of the arcs (possibly both). We denote by $P_{u,v}$ an arbitrary directed shortest path from u to v . For $i < j$ we denote by $S[i, j]$ a path (along the infinite face) $s_i - s_{i+1} - \dots - s_j$. We extend the definition of $rev(\cdot)$ to paths. We denote by $P[w, y]$ the subpath of P between vertices w and y . We similarly use $P(w, y]$, $P[w, y)$, and $P(w, y)$ to denote whether the subpath includes the corresponding endpoint(s) or not. We use \circ to denote a concatenation of two paths. Let C be a directed non-crossing cycle in G . We denote by $left(C)$ and $right(C)$ the subgraphs of G that consist of all edges, vertices and faces that are lying to the left and right of C , respectively. The arcs of C and their reverses are in both $left(C)$ and $right(C)$.




Figure 2: An example of two directed paths P and Q that cross at the crossing part R .

Let P and Q be directed paths or cycles. We say that they *cross at subpath R* if, when ignoring their orientation: (1) R is a proper (not a prefix or suffix) subpath of both P and Q , and (2) The edges of Q that follow and precede R are in different sides of P . See Figure 2. We refer to R as a *crossing part* of P and Q .

The dual graph of G is denoted by $G^* = (V^*, E^*)$. Again, we think of G^* as a directed graph with a set of arcs \mathcal{A}^* , defined as follows. For every arc $uv \in \mathcal{A}$, there is a corresponding arc $(uv)^* \in \mathcal{A}^*$ such that the tail and head of $(uv)^*$ are the faces that lie to the right and the left of uv , respectively. We note that we slightly abuse the notation here, since the dual of $(uv)^*$ is $rev(uv)$ (and not uv). For $B \subseteq \mathcal{A}$, let $B^* = \{(uv)^* \mid uv \in B\}$. For a cut $X \subseteq V$, let $\delta(X) = \{uv \in \mathcal{A} \mid u \in X, v \in V \setminus X\}$. For a cycle C^* in the dual graph we say that $v \in V$ is in $left(C^*)$ (resp. $right(C^*)$) if the face of G^* that corresponds to v is in $left(C^*)$ (resp. $right(C^*)$).

3 A Bisector-Based Approach to the Okamura-Seymour Compression

In this section we present our new proof of Theorem 1 and the proofs of Lemma 2 and Theorem 3. Our main tool is the use of simple dual cycles that we call *bisectors*. In Section 3.1 we define bisectors, and prove that they are arc-disjoint and that this implies Lemma 2. In Section 3.2 we use it to prove Theorem 3. Then, in Section 3.3 we show that the union of all bisectors partitions the graph into regions such that all vertices belonging to the same region have the same pattern. Finally, in Section 3.4 we show that every two bisectors can cross at most $O(k)$ times, implying that the partition induces only $O(k^3)$ regions (and hence only $O(k^3)$ distinct patterns) thus proving Theorem 1.

3.1 Bisectors

For $1 \leq i \leq k-1$, define the cut $A_i = \{v \in V \mid p_v[i] = -1\}$. Since we assume the patterns are over $\{-1, 1\}$, $V \setminus A_i = \{v \in V \mid p_v[i] = 1\}$. We define the *bisector* $\beta_i = \delta(A_i)^*$. Namely, β_i consists of all arcs $(uv)^* \in \mathcal{A}^*$ such that $p_u[i] = -1$ and $p_v[i] = 1$. Moreover, every edge $\{u, v\} \in E$ such that $p_u \neq p_v$ belongs to some bisector (possibly more than one). By cut-cycle duality, if the induced subgraphs of A_i and $V \setminus A_i$ are both connected, then β_i is a directed simple cycle in the dual graph. The next lemma implies that both induced subgraphs of A_i and $V \setminus A_i$ are connected.




Figure 3: The bisector β_i and its corresponding cut A_i .

Lemma 5. *For any $u \in A_i$ (resp. $V \setminus A_i$), the vertices of $P_{u, s_{i+1}}$ (resp. P_{u, s_i}) are in A_i (resp. $V \setminus A_i$).*

Proof. Assume that $u \in A_i$ (the proof of the other case is symmetric). Let v be any vertex of $P_{u, s_{i+1}}$ and assume for the sake of contradiction that $v \in V \setminus A_i$. Thus, $d(v, s_{i+1}) = d(v, s_i) + 1$. By the triangle inequality, we get the following contradiction:

$$d(u, s_i) = d(u, s_{i+1}) + 1 = d(u, v) + d(v, s_{i+1}) + 1 = d(u, v) + d(v, s_i) + 2 \geq d(u, s_i) + 2 \quad \square$$

By the above lemma, any two vertices $u, v \in A_i$ (resp. $V \setminus A_i$) are connected in the induced subgraph of A_i (resp. $V \setminus A_i$) by the path $P_{u, s_{i+1}} \circ \text{rev}(P_{v, s_{i+1}})$ (resp. $P_{u, s_i} \circ \text{rev}(P_{v, s_i})$). This yields the following corollary.

Corollary 6. β_i is a directed simple cycle in the dual graph.

We note that the above corollary implies that for any face f , every bisector contains at most two arcs incident to f . This shows that there are only $O(k)$ total bit changes between patterns as we go along the vertices of a face f .

Another useful corollary comes from the fact that any edge whose dual is in β_i contains endpoints that are both in A_i and $V \setminus A_i$. Therefore:

Corollary 7. For any $u \in A_i$ (resp. $u \in V \setminus A_i$), the dual edges of $P_{u,s_{i+1}}$ (resp. P_{u,s_i}) are not in β_i .

Note that β_i has two arcs incident to f_∞ , one of them being $(s_{i+1}s_i)^*$. We think of $(s_{i+1}s_i)^*$ as the first arc of β_i . See Figure 3. The following lemma shows that bisectors are arc-disjoint.

Lemma 8. Every pair of bisectors β_i, β_j are arc-disjoint.

Proof. Assume for contradiction that arc $(uv)^*$ appears both in β_i and in β_j . By definition, u belongs to A_i and A_j , and v belongs to $V \setminus A_i$ and $V \setminus A_j$. We first prove that under our assumption, either $P_{u,s_{i+1}}$ intersects with P_{v,s_j} or $P_{u,s_{j+1}}$ intersects with P_{v,s_i} . To see why, first note that since P_{v,s_i} and P_{v,s_j} are shortest paths, we can choose them to follow a common maximal-length prefix $P_{v,s_i}[v, w] = P_{v,s_j}[v, w]$ for some w , and they do not intersect again after w . Consider the directed cycle $C = rev(P_{v,s_j}) \circ P_{v,s_i} \circ S[i, j]$ (see Figure 4). Notice that by our choice of P_{v,s_i} and P_{v,s_j} and by the fact that $S[i, j]$ lies on the infinite face, C is not necessarily simple but it does not self-cross. We have two cases to consider:

Case 1: $u \in \text{left}(C) \setminus C$. Since $s_{i+1} \in \text{right}(C)$ then (by the Jordan curve theorem and the fact that all vertices of S lie on the infinite face) $P_{u,s_{i+1}}$ must intersect with $rev(P_{v,s_j}) \circ P_{v,s_i}$. However, by Lemma 5, $P_{u,s_{i+1}}$ cannot intersect with P_{v,s_i} , therefore it intersects with $rev(P_{v,s_j})$ (and hence with P_{v,s_j}).

Case 2: $u \in \text{right}(C)$. Notice that $s_{j+1} \in \text{left}(C)$. If s_{j+1} is in P_{v,s_i} then $P_{u,s_{j+1}}$ intersects with P_{v,s_i} and we are done. Otherwise, since s_{j+1} is not in P_{v,s_j} by Lemma 5, then $s_{j+1} \in \text{left}(C) \setminus C$. But then again (by the Jordan curve theorem) $rev(P_{v,s_j})$ (and hence $P_{u,s_{j+1}}$) must intersect with $rev(P_{v,s_j}) \circ P_{v,s_i}$. However, by Lemma 5, $P_{u,s_{j+1}}$ cannot intersect with $rev(P_{v,s_j})$, therefore it intersects with P_{v,s_i} .

We can therefore continue under the assumption that $P_{u,s_{i+1}}$ and P_{v,s_j} intersect at a vertex x (the other case is symmetric). By the triangle inequality:

$$\begin{aligned} d(u, s_{j+1}) &\leq d(u, x) + d(x, s_j) - 1 \\ d(u, x) + d(x, s_{i+1}) &\leq d(u, v) + d(v, s_i) - 1 \\ d(v, s_i) &\leq d(v, x) + d(x, s_{i+1}) - 1 \\ d(v, x) + d(x, s_j) &\leq d(v, u) + d(u, s_{j+1}) - 1 \end{aligned}$$

Since $d(u, v) = d(v, u) = 1$, summing the above inequalities we get the contradiction $0 \leq -2$. \square

The above lemma shows that two bisectors cannot share an arc. Note however that it is still possible that a bisector contains *reversed* arcs of another bisector. This proves Lemma 2. Next, we prove Theorem 3.




Figure 4: The two cases in the proof of Lemma 8. Case 1 on the left. Case 2 on the right. $\text{right}(C)$ is shaded. For clarity, C is a simple cycle in this example.

3.2 A proof of Theorem 3

We begin by describing how to compute all the bisectors of the graph and report their arcs in $\tilde{O}(n)$ time. We split every edge $\{s_i, s_{i+1}\}$ by adding a dummy vertex y_i and edges $\{s_i, y_i\}, \{y_i, s_{i+1}\}$ of weight $\frac{1}{2}$. Consider a shortest path tree T_i rooted at y_i . Notice that the arcs of β_i which are not incident to f_∞ , are the duals of arcs whose tail is in the subtree rooted at s_{i+1} and head is in the subtree rooted at s_i . In the interdigitating tree of T_i (i.e., the tree in the dual graph whose edges are the duals of the edges not in T_i), they are precisely the f_i -to- f_∞ path without the last arc, where $f_i \neq f_\infty$ is the face incident to $\{s_i, s_{i+1}\}$ in G . We can therefore run the MSSP algorithm of Klein [22] in $\tilde{O}(n)$ time, and report for every $1 \leq i \leq k-1$ all the those arcs of β_i in $\tilde{O}(|\beta_i|)$ time. To report the two arcs of β_i which are incident to f_∞ , one of them is trivially $(s_{i+1}s_i)^*$ and the other one is determined by last arc of the above f_i -to- f_∞ path. Since by Lemma 8 the arcs of bisectors are disjoint, this takes time $\tilde{O}(n + \sum_{i=1}^{k-1} |\beta_i|) = \tilde{O}(n)$ time. In particular, we can label every edge $\{u, v\} \in E$ by the (at most two) bisectors that use $(uv)^*$ and $(vu)^*$. I.e., the bits that change between p_u and p_v .

We next describe the compression scheme. Recall that, by storing $d(v, s_1)$ for every $v \in T$, a query $d(v, s_i)$ (with $v \in T$ and $s_i \in S$) boils down to extracting p_v and computing its $(i-1)$ 'th prefix-sum. Let \mathcal{T} be a spanning tree of G . Label each edge $\{u, v\}$ of \mathcal{T} by the (at most two) bits that change between the patterns of u and v . Note that there could be many (potentially $\Omega(n)$) nodes of \mathcal{T} that correspond to the same pattern. In order to decrease the size of \mathcal{T} to be $p_\#$ (the number of distinct patterns in G), we root \mathcal{T} at some arbitrary node u . Then, for every two nodes v, w of \mathcal{T} s.t. $p_v = p_w$ (and w.l.o.g. v is not a descendent of w) we remove the node w and turn all its children to be children of v (their edge labels remain the same). We repeat this process until the size of the tree is $p_\#$. We denote the resulting tree by \mathcal{T}' . Let Q be an Euler-tour of \mathcal{T}' starting from the root u . Consider the patterns of the nodes as we go along Q , starting from p_u . In each step, the pattern only changes in at most two bits (according to the edge labels). Therefore, we can maintain all these $O(|Q|) = O(p_\#)$ versions of the pattern using a *persistent* [11] data structure for prefix-sum (e.g., using persistent segment trees [3]). Such a data structure supports both updates and prefix-sum queries to any version in $\tilde{O}(1)$ time and uses $\tilde{O}(|Q| + k) = \tilde{O}(p_\# + k)$ space. Finally,

for every vertex $v \in T$ let q_v be a node in Q whose corresponding pattern is p_v . We store a pointer from v to the version of the persistent data structure at q_v , using additional $\tilde{O}(|T|)$ bits overall.

We now give a randomized $\tilde{O}(n)$ time algorithm for constructing \mathcal{T}' (and hence the compression). An arbitrary spanning tree \mathcal{T} can be computed in $O(n)$ time. Assume that every edge $\{w, v\}$ of \mathcal{T} is labeled by the (at most two) bits that change between p_w and p_v . Let us compute the pattern of the root u of \mathcal{T} with a single-source shortest-paths computation in G . We also compute the *Karp-Rabin fingerprint* [20] $\phi(p_u)$ of p_u . Such fingerprints are appealing because: (1) for any $p_w \neq p_v$, we have that $\phi(p_w) \neq \phi(p_v)$ with high probability, and (2) given $\phi(S_1)$ and $\phi(S_2)$ of two strings S_1, S_2 we can compute in $O(1)$ time the fingerprint of the concatenation $\phi(S_1 \circ S_2)$. Thus, if we maintain a complete binary tree on top of the pattern where each node contains the fingerprint of its subtree (and in particular, the root contains the fingerprint of the entire pattern), then we can update this tree in $O(\log k)$ time after changing one or two bits in the pattern.

We maintain the fingerprints in a dictionary initially containing only $\phi(p_u)$. We process the nodes of \mathcal{T} starting from u , maintaining a queue of next-to-visit nodes. When we process a node v , we compute $\phi(p_v)$ from the fingerprint of v 's parent, by flipping the bits according to the edge label (in $O(\log k)$ time). We then try to add $\phi(p_v)$ to the dictionary. If we find a collision with some vertex w (namely, $\phi(p_v) = \phi(p_w)$) then we delete v from \mathcal{T} , and set the children of v to be children of w in \mathcal{T} . In any case, we add the children of v to the queue so they will be processed later. Notice that a node is visited only after all its ancestors have been visited. Therefore, we can always compute its fingerprint and we never move children from a vertex to its descendent, so \mathcal{T} remains a tree. In addition, the parent of every node changes or gets deleted at most once, hence the running time is $\tilde{O}(n)$. Overall, in $\tilde{O}(n)$ time we construct \mathcal{T}' and the dictionary (both of size $p\#$).

- In the special case where the vertices of T induce a connected component in G , we can skip the first part of the algorithm and simply take a path Q that traverses only the vertices of T . The rest of the construction remains the same and since $|Q| = O(|T|)$, the size of the compression is $\tilde{O}(|T| + k)$.
- In the special case where the vertices of T all lie on a single face (but not necessarily consecutively), let Q be a path that visits all the vertices of the face in clockwise order. By Corollary 6, the total number of bit changes between patterns of consecutive vertices along Q is $O(k)$. Therefore, the number of patterns encountered is $O(k)$ and hence we get an $\tilde{O}(|T| + k)$ compression for this case as well.

This completes the proof of Theorem 3.

3.3 The bisector graph and the pattern graph

The *bisector graph* $G_{\mathcal{B}}$ is the subgraph of G^* composed of the union of all the bisectors. The faces of $G_{\mathcal{B}}$ represent the patterns of G in the following way.

Lemma 9. *For every $u, v \in V$, if u and v are embedded inside the same face f of $G_{\mathcal{B}}$, then $p_u = p_v$.*

Proof. Notice that $G_{\mathcal{B}}$ is a connected graph because all the bisectors are incident to f_{∞} . Hence, f is a simple cycle in G^* . Let G_f be the subgraph of G embedded inside the face f . Since there are no bisector edges embedded inside f , then in G_f there is no pair of adjacent vertices that have

different patterns. Since f is a simple cycle, then by cut-cycle duality G_f is a connected subgraph. Therefore, there exists a u -to- v path in G_f , and every pair of adjacent vertices in this path have the same pattern. Hence $p_u = p_v$. \square

By the above lemma, every pattern p of G corresponds to a unique nonempty subset of faces of $G_{\mathcal{B}}$. More precisely, a pattern p corresponds to all the faces of $G_{\mathcal{B}}$ such that the vertices of G embedded in these faces have pattern p . In particular, the number of faces of $G_{\mathcal{B}}$ is an upper bound on the number of distinct patterns in G . Therefore, if we could prove that $G_{\mathcal{B}}$ has $O(k^3)$ faces we would be done. Unfortunately, this is not the case. There can be as many as $\Omega(n)$ faces of $G_{\mathcal{B}}$ that correspond to the same pattern (see Figure 5). To tackle this, we transform $G_{\mathcal{B}}$ into a new graph $G_{\mathcal{P}}$ (called the *pattern graph*) that has only $O(k^3)$ faces and whose faces still represent all the distinct patterns of G .




Figure 5: The shaded faces all correspond to the same pattern (assuming no other bisector crosses or separates them). They are formed when the two bisectors either cross each other or just touch (intersect without crossing). We will later see that two bisectors can cross each other at most $O(k)$ times, but, they can touch $\Omega(n)$ times, creating $\Omega(n)$ faces that correspond to the same pattern.

The *pattern graph* $G_{\mathcal{P}}$ is obtained by applying on $G_{\mathcal{B}}$ the following two-phase procedure:

- (1) A *Peel* phase: Recall that while Lemma 8 says that every two bisectors are arc-disjoint, it is still possible that one bisector contains *reversed* arcs of another. In the peel phase, we re-embed the bisectors so that no bisector contains reversed arcs of another bisector. After the peel phase, crossings and touchings occur only at vertices (rather than subpaths).
- (2) A *Merge* phase: In the merge phase, we merge faces that correspond to the same pattern and share a common vertex.

Peel Phase. For every two bisectors β and β' , consider the set of maximal-length subpaths R , such that R is a subpath of β and $rev(R)$ is a subpath of β' . If the arc of β that follows R is in $right(\beta')$ (resp. $left(\beta')$), then we re-embed every arc of R on a new curve lying to the right (resp. left) of its reverse. See Figure 6. Note that the peel phase does not create any new crossings between β and β' .




Figure 6: Before (left) and after (right) a peel phase. In this example, R is a single arc, and the arc of β that follows R is in $\text{left}(\beta')$.

Merge Phase. For every vertex $g \neq f_\infty$ of a bisector β , if any other bisector crosses β at g then we do nothing. Otherwise, we split g into two copies. All the arcs in $\text{left}(\beta)$ that are incident to g are connected to one copy, and all the arcs in $\text{right}(\beta)$ that are incident to g are connected to the other copy. Finally, we replace the arcs of β that are incident to g (say fg and gh) by a single arc fh . See Figure 7. Note that if g is not incident to any bisector other than β , then the merge phase simply contracts the arc gh . We repeat this process until there are no such bisector pairs in the graph.




Figure 7: Before (left) and after (right) a merge phase. The (shaded) face g' is obtained by merging faces h' and f' .

We now show that the above two-phase procedure maintains the relation between patterns in G and faces in $G_{\mathcal{P}}$. Namely, that every pattern in G corresponds to a unique nonempty subset of faces of $G_{\mathcal{P}}$. To this end, we extend the definition of patterns to faces of $G_{\mathcal{B}}$. This step is necessary since the peel phase creates faces that do not correspond to primal vertices.

We define the *pattern* of a face f of $G_{\mathcal{B}}$, denoted p_f , to be the length $k - 1$ vector where $p_f[i] = -1$ (resp. $p_f[i] = 1$) if f is a face in $\text{left}(\beta_{i+1})$ (resp. $\text{right}(\beta_{i+1})$) in $G_{\mathcal{B}}$. The definition remains the same for any graph we obtain from $G_{\mathcal{B}}$ during the two-phase procedure. The following two propositions show that this definition is consistent with the original definition of patterns (of vertices).

Proposition 10. *Let $v \in V$ be a vertex embedded inside a face f of $G_{\mathcal{B}}$. Then $p_v = p_f$.*

Proof. Let $0 \leq i \leq k - 2$. If $p_f[i] = 1$ then f is in $\text{right}(\beta_{i+1})$ by definition. Since $G_{\mathcal{B}}$ is a subgraph of G^* , then in G^* v is also embedded in $\text{right}(\beta_{i+1})$. Hence, $v \in V \setminus A_i$ and therefore $p_v[i] = 1$. A symmetric argument shows that if $p_f[i] = -1$ then $p_v[i] = -1$. \square

By Proposition 10, all the faces of $G_{\mathcal{B}}$ that correspond to a pattern p_v have the same face pattern. Notice that the peel phase does not change the patterns of existing faces. It can only

add new faces to the graph, but no vertex of G is embedded in any of these new faces. Hence, the relation is preserved after the peel phase. Next we show that after a merge step, every pattern still corresponds to a unique subset of faces (i.e., we show that we do not merge faces that corresponded to different patterns). Consider a single merge step happening at g (as illustrated in Figure 7). Denote by f' (resp. h') the face lying to the left of fg (resp. gh). Namely, f' and h' are the faces that get merged (a symmetric argument holds when they lie to the right of fg and gh). Let g' denote the face obtained by merging f' and h' .

Proposition 11. $p_{f'} = p_{h'} = p_{g'}$.

Proof. Since no bisector crosses β at g , then fg and gh belong to the same side of every bisector. This, together with the fact that fg , gh , and their reverses do not belong to any other bisector, implies that f' and h' also belong to the same side of every bisector. Hence $p_{f'} = p_{h'}$. Now consider the arc fh after the merge. Since f and h belong to the same side of every bisector as fg (and gh), then g' also belongs to the same side of every bisector, hence $p_{f'} = p_{h'} = p_{g'}$. \square

By proposition 10, if f' or h' are faces that correspond to p_v then they do not correspond to any $p_u \neq p_v$. By Proposition 11, we can set g' to correspond to p_v , and the set of faces corresponding to every pattern remains unique. This yields the following corollary.

Corollary 12. *Every pattern of G corresponds to a unique subset of faces of $G_{\mathcal{P}}$.*

Finally, we show that the number of faces in $G_{\mathcal{P}}$ depends linearly on the number of bisector crossings. Let t be the total number of bisector crossings in $G_{\mathcal{P}}$. That is, t is the sum of the number of crossings between all pairs of bisectors.

Lemma 13. *The number of faces in $G_{\mathcal{P}}$ is $O(t + k)$.*

Proof. By Euler's formula, it suffices to show that the number of arcs in $G_{\mathcal{P}}$ is $O(t + k)$. For every arc fg in $G_{\mathcal{P}}$, where neither f nor g is f_{∞} , the arc fg belongs to some bisector β . Moreover, there must exist some other bisector that crosses β at f . Otherwise, the arc would have been removed in the merge phase. Consider all the bisectors that cross β at f in a clockwise order around f starting at fg . Let β' be the one following fg . Then we charge fg to the crossing of β and β' at f . Notice that at most two arcs will be charged to this crossing of β and β' (the arc fg and the arc of β' whose tail is f). Overall, we have charged $O(t)$ arcs. The only arcs that did not get charged are the $2(k - 1)$ arcs incident to f_{∞} . Therefore, the number of arcs in $G_{\mathcal{P}}$ is $O(t + k)$. \square

In the next subsection, we prove that every pair of bisectors can cross at most $O(k)$ times. Since there are $O(k^2)$ pairs of bisectors, the total number of crossings is then $t = O(k^3)$, which by Corollary 12 and Lemma 13 implies Theorem 1.

3.4 Two bisectors can cross only $O(k)$ times

Let β_i and β_j be two bisectors in G^* that cross each other at least once. Let R_1, R_2, \dots, R_r be their crossing parts that do not contain f_{∞} , sorted by their order of appearance along β_i . We note that since β_i and β_j are simple cycles, the crossing parts must be disjoint. In Lemma 16 we show that the crossing parts appear in reverse order along β_j , and in Lemma 17 we use this fact to prove that the number of crossings r is at most $O(k)$. We begin by defining an important configuration of bisectors and shortest paths.




Figure 8: A configuration of bisectors β_i, β_j and shortest paths $P_{u, s_{i+1}}, P_{v, s_j}$ that must intersect. $\text{right}(C^*)$ is shaded.

Lemma 14. *Let $C^* = \beta_j[f_\infty, f] \circ \text{rev}(\beta_i[f_\infty, f])$ be a simple cycle, and let $u \in A_i, v \in V \setminus A_j$ (resp. $u \in V \setminus A_i, v \in A_j$) be two vertices in $\text{left}(C^*)$ (resp. $\text{right}(C^*)$). Then $P_{u, s_{i+1}}$ and P_{v, s_j} (resp. P_{u, s_i} and $P_{v, s_{j+1}}$) must intersect at some vertex x .*

Proof. We focus on the case where $u \in A_i$ and $v \in V \setminus A_j$ are vertices in $\text{left}(C^*)$ (the proof of the other case is symmetric). See Figure 8. We assume that G and G^* are embedded on the same surface, such that for every $wy \in \mathcal{A}$, the curves of wy and $(wy)^*$ intersect in their middles at a single point p on the surface. See Figure 9. We refer to the two parts of the curve of wy as $(w - p)$




Figure 9: An arc $wy \in \mathcal{A}$ that intersects with its dual $(wy)^* \in \mathcal{A}^*$ at a middle point p .

and $(p - y)$. For a path Q that contains arc wy , we slightly abuse notation and use $Q[\cdot, p]$ and $Q[p, \cdot]$ to denote a prefix and suffix of the curve of Q . In addition, we say that a path P of the primal graph crosses a path Q of the dual graph, if P contains an arc wy whose dual or reversed dual is in Q . In particular, it means that there exist a common point p (in the middle of wy), such that $(w - p)$ and $(p - y)$ are on different sides of Q . Let p_j be the point in the middle of $s_j s_{j+1}$, and let p_i be the point in the middle of $s_i s_{i+1}$.

Notice that $s_{i+1}, s_j \in \text{right}(C^*) \setminus C^*$ by definition, and that $v \in \text{left}(C^*) \setminus C^*$ by assumption (and the fact that v is not part of C^* because v is a primal vertex). Hence, P_{v, s_j} must cross C^* . However, by Corollary 7 it cannot cross $\beta_j[f_\infty, f]$, hence it must cross $\beta_i[f_\infty, f]$. This means that there is a point q that is common to both P_{v, s_j} and $\beta_i[f_\infty, f]$. In particular, let q be the last point along the curve of P_{v, s_j} that is also along the curve of C^* . Notice that $P_{v, s_j}[q, s_j] \circ (s_j - p_j)$ is a chord inside the cycle C^* . Similarly, there exists a point q' along the curve of $P_{u, s_{i+1}}$ such that $P_{u, s_{i+1}}[q', s_{i+1}] \circ (s_{i+1} - p_i)$ is a chord in C^* . The endpoints of the chords appear in clockwise order

along C^* as (q, p_i, p_j, q') . It is well known that two chords in such a configuration must intersect. Therefore, there exist a primal vertex x that is common to both $P_{v, s_j}[q, s_j]$ and $P_{u, s_{i+1}}[q', s_{i+1}]$. \square

It is important to remark that Lemma 14 holds even when the cycle C^* is a non-self-crossing non-simple cycle. Namely, if C^* intersects with itself, we let g be the first intersection vertex in $\beta_i[f_\infty, f]$, and define a cycle $\hat{C}^* = \beta_j[f_\infty, g] \circ \text{rev}(\beta_i[f_\infty, g])$. Note that since β_i and β_j are simple cycles, and by our choice of g , there are no intersections in \hat{C}^* . Clearly, we also have that $v, u \in \text{left}(\hat{C}^*)$. Thus, we apply the Lemma to \hat{C}^* instead of C^* .

Corollary 15. $P_{v, s_j}[v, x]$ (resp. $P_{u, s_{i+1}}[u, x]$) contains an edge whose dual is in β_i (resp. β_j).

We are now in the position to prove the two main lemmas of this section.

Lemma 16. Let β_i and β_j be two crossing bisectors. Let R_1, R_2, \dots, R_r be their crossing parts along β_i . Then, the crossing parts along β_j are reversed R_r, R_{r-1}, \dots, R_1 .

Proof. We assume that $i < j$. For $\ell \leq r$, we say that β_j enters R_ℓ from $\text{left}(\beta_i)$ (resp. $\text{right}(\beta_i)$) if the arc of β_j that precedes R_ℓ is in $\text{left}(\beta_i)$ (resp. $\text{right}(\beta_i)$). For brevity, we assume that every R_ℓ is a single vertex in V^* . Finally, we assume without loss of generality that s_j is in $\text{left}(\beta_i)$ (the proof of the other case is symmetric).

Assume for the sake of contradiction that the order of appearance is not the reverse order. Then there exists a pair of crossing parts R_a and R_b such that: (1) $a < b$, (2) R_b is the crossing part following R_a in β_j , and (3) a is minimal among such pairs. Consider the cycle $C^* = \beta_j[f_\infty, R_a] \circ \text{rev}(\beta_i[f_\infty, R_a])$. Note that C^* is non-crossing since if there exists R_c for $c < a$ that $\beta_j[f_\infty, R_a]$ crosses, then a wouldn't be minimal. We have two cases to consider:




Figure 10: Case 1 in the proof of Lemma 16. The dotted parts represent parts in which there may be crossings. $\text{right}(C^*)$ is shaded. In this example, P is crossed by $\beta_j[f_\infty, R_a]$ exactly twice.

Case 1: β_j enters R_a from $\text{left}(\beta_i)$ (hence it enters R_b from $\text{right}(\beta_i)$). Let $(uv)^*$ be the arc of β_j that follows R_b . We next show that v and C^* form the configuration of Lemma 14, in the special case of $v = u$.

First we show that $v \in V \setminus A_j$ and $v \in A_i$. Note that $v \in V \setminus A_j$ by definition, since v is the primal vertex lying to the right of an arc of β_j . Since $(uv)^*$ follows the crossing part R_b , and since β_j enters R_b from $\text{right}(\beta_i)$, then $(uv)^* \in \text{left}(\beta_i) \setminus \beta_i$. Hence, $v \in \text{left}(\beta_i)$ and therefore $v \in A_i$.

Next we show that $v \in \text{left}(C^*)$. For this, we show that $R_b \in \text{left}(C^*) \setminus C^*$, hence the arcs incident to R_b and their corresponding primal vertices are in $\text{left}(C^*)$. Consider the path $P = \beta_i[R_a, R_b]$. The first arc of P is in $\text{left}(C^*)$ by the case assumption. To show that $R_b \in \text{left}(C^*) \setminus C^*$ we will show that P crosses the cycle C^* an even number of times. Notice that P does not cross $\beta_i[f_\infty, R_a]$ since β_i is a simple cycle. It therefore remains to show that P crosses $\beta_j[f_\infty, R_a]$ an even number of times (equivalently, we show that $\beta_j[f_\infty, R_a]$ crosses P an even number of times). To this end, let us define a cycle $\hat{C}^* = P \circ \text{rev}(\beta_j[R_b, R_a])$. Notice that \hat{C}^* is non-crossing since R_b follows R_a in β_j . Notice that the first and last arcs of $\beta_j[f_\infty, R_a]$ are both in $\text{left}(\beta_i)$. Also notice that $\text{right}(\hat{C}^*)$ is included in $\text{right}(\beta_i)$ (by the case assumption), therefore the first and last arcs of $\beta_j[f_\infty, R_a]$ are in $\text{left}(\hat{C}^*)$. Hence, $\beta_j[f_\infty, R_a]$ crosses \hat{C}^* an even number of times. Since $\beta_j[f_\infty, R_a]$ can only cross \hat{C}^* at P (as β_j is a simple cycle), it must cross P an even number of times.

We can thus apply Lemma 14, and conclude that P_{v,s_j} and $P_{v,s_{i+1}}$ intersect at vertex x , and that $P_{v,s_{i+1}}[v, x]$ contains an edge whose dual is in β_j (by Corollary 15). Since the lengths of $P_{v,s_j}[v, x]$ and $P_{v,s_{i+1}}[v, x]$ are the same, $P_{v,s_{i+1}}[v, x] \circ P_{v,s_j}[x, s_j]$ is a shortest v -to- s_j path. However, since $P_{v,s_{i+1}}[v, x]$ contains an edge of β_j , we get a contradiction to Corollary 7.




Figure 11: Case 2 in the proof of Lemma 16. The dotted parts represent parts in which there may be crossings. $\text{left}(C^*)$ is shaded. In this example, P is crossed by $\beta_j[f_\infty, R_a]$ exactly twice.

Case 2: β_j enters R_a from $\text{right}(\beta_i)$ (hence it enters R_b from $\text{left}(\beta_i)$). Let $(uv)^*$ be the arc of β_j that follows R_b . We next show that u and C^* form the configuration of Lemma 14 (again, in the special case of $u = v$). By symmetric arguments to Case 1, we can show that $u \in A_j$ and $u \in V \setminus A_i$. We therefore only need to show that $u \in \text{right}(C^*)$.

To show that $u \in \text{right}(C^*)$, it suffices to show that $R_b \in \text{right}(C^*) \setminus C^*$ (hence the arcs incident to R_b and their corresponding primal vertices are in $\text{right}(C^*)$). Consider the path $P = \beta_i[R_a, R_b]$. The first arc of P is in $\text{right}(C^*)$ by the case assumption. To show that $R_b \in \text{right}(C^*) \setminus C^*$, we will show that P crosses the cycle C^* an even number of times. Since P does not cross $\beta_i[f_\infty, R_a]$, we need to show that $\beta_j[f_\infty, R_a]$ crosses P an even number of times (here is where the argument will differ from Case 1).

Let $\hat{C}^* = P \circ \text{rev}(\beta_j[R_b, R_a])$ be a non-crossing cycle. Let R be the first crossing between $\beta_j[f_\infty, R_a]$ and β_i . R exists and is not R_a by the assumption that $s_j \in \text{left}(\beta_i)$, and by the case assumption. Note that $\beta_j[f_\infty, R]$ does not cross β_i (and hence does not cross P) by definition of R ,

so it remains to show that the number of crossings between $\beta_j[R, R_a]$ and P is even. Notice that, since $s_j \in \text{left}(\beta_i)$, the first and last arcs of $\beta_j[R, R_a]$ are both in $\text{right}(\beta_i)$. Also note that $\text{left}(\hat{C}^*)$ is contained in $\text{left}(\beta_i)$. Therefore, the first and last arcs of $\beta_j[R, R_a]$ are in $\text{right}(\hat{C}^*)$. Hence, $\beta_j[R, R_a]$ crosses \hat{C}^* an even number of times. Since it can only cross \hat{C}^* at P , then $\beta_j[R, R_a]$ crosses P an even number of times.

We can thus apply Lemma 14, and conclude that $P_{u, s_{j+1}}$ and P_{u, s_i} intersect at vertex x , and that $P_{u, s_{j+1}}[u, x]$ contains an edge of β_i (by Corollary 15). Since the lengths of $P_{u, s_i}[u, x]$ and $P_{u, s_{j+1}}[u, x]$ are the same, $P_{u, s_{j+1}}[u, x] \circ P_{u, s_i}[x, s_i]$ is a shortest u -to- s_i path. However, since $P_{u, s_{j+1}}[u, x]$ contains an edge of β_i , we get a contradiction to Corollary 7. \square

Lemma 17. *Two bisectors can cross at most $\frac{k}{2} + O(1)$ times.*

Proof. We will prove that if two bisectors β_i, β_j cross r times then there exists a vertex $v \in V$ such that $d(v, s_j) - d(v, s_{i+1}) \geq 2r - \frac{k}{2} - O(1)$. Since the distance between any pair of vertices along the infinite face is at most $\frac{k}{2}$, then by the triangle inequality we have also $d(v, s_j) - d(v, s_{i+1}) \leq \frac{k}{2}$. Hence, we get $2r - \frac{k}{2} - O(1) \leq \frac{k}{2}$ and the lemma follows.




Figure 12: Two bisectors that cross r times and the vertices and shortest paths they induce. Observe that v_ℓ is in $\text{right}(\beta_i)$ when $r - \ell$ is even, and otherwise in $\text{left}(\beta_i)$. $\text{right}(C_\ell^*)$ is shaded.

Again, let us assume without loss of generality that s_j is in $\text{left}(\beta_i)$ (the proof of the other case is symmetric). For every $\ell < r$, consider the cycle $C_\ell^* = \beta_j[f_\infty, R_\ell] \circ \text{rev}(\beta_i[f_\infty, R_\ell])$. By the previous lemma, C_ℓ^* does not self-cross. For even (resp. odd) $r - \ell$, let $(u_\ell v_\ell)^*$ (resp. $(v_\ell u_\ell)^*$) be the arc of β_j that follows R_ℓ . See Figure 12. We assume that $r - \ell$ is even (the odd case is symmetric).

We claim that $C_\ell^*, v_\ell, v_{\ell+1}$ forms the configuration of Lemma 14. By definition of β_j we have that $v_\ell \in V \setminus A_j$. Thus, it remains to prove that $v_{\ell+1} \in A_i$ and that $v_\ell, v_{\ell+1} \in \text{left}(C_\ell^*)$. By Lemma 16 and the assumption that $s_j \in \text{left}(\beta_i)$, β_j reaches $R_{\ell+1}$ from $\text{right}(\beta_i)$. Therefore, $(v_{\ell+1} u_{\ell+1})^* \in \text{left}(\beta_i) \setminus \beta_i$. Hence, $v_{\ell+1}$ is in $\text{left}(\beta_i)$ and therefore $v_{\ell+1} \in A_i$. To see that $v_\ell, v_{\ell+1} \in \text{left}(C_\ell^*)$, note that by Lemma 16, $(v_{\ell+1} u_{\ell+1})^*$ is in $\beta_j[f_\infty, R_\ell]$, hence $v_{\ell+1} \in \text{left}(C_\ell^*)$ by definition (as $\beta_j[f_\infty, R_\ell]$

is part of C_ℓ^*). Clearly, since β_j enters R_ℓ from $\text{right}(\beta_i)$ we also have that $(u_\ell v_\ell)^* \in \text{left}(C_\ell^*)$ and in particular $v_\ell \in \text{left}(C_\ell^*)$.

We can thus apply Lemma 14 and conclude that P_{v_ℓ, s_j} and $P_{v_{\ell+1}, s_{i+1}}$ must intersect at some vertex x_ℓ . In addition, we can conclude that $P_{v_{\ell-1}, s_{j+1}}$ and P_{v_ℓ, s_i} intersect at some vertex $x_{\ell-1}$, and that $P_{v_{\ell-2}, s_j}$ and $P_{v_{\ell-1}, s_{i+1}}$ intersect at some vertex $x_{\ell-2}$. Therefore, by the triangle inequality and the assumption that patterns are binary, we have:

$$\begin{aligned} d(v_\ell, x_\ell) + d(x_\ell, s_j) &\leq d(v_\ell, x_{\ell-1}) + d(x_{\ell-1}, s_{j+1}) - 1 \\ d(v_\ell, x_{\ell-1}) + d(x_{\ell-1}, s_i) &\leq d(v_\ell, x_\ell) + d(x_\ell, s_{i+1}) - 1 \\ d(v_{\ell-1}, x_{\ell-2}) + d(x_{\ell-2}, s_{i+1}) &\leq d(v_{\ell-1}, x_{\ell-1}) + d(x_{\ell-1}, s_i) - 1 \\ d(v_{\ell-1}, x_{\ell-1}) + d(x_{\ell-1}, s_{j+1}) &\leq d(v_{\ell-1}, x_{\ell-2}) + d(x_{\ell-2}, s_j) - 1 \end{aligned}$$

Summing the above inequalities we get:

$$d(x_\ell, s_j) - d(x_\ell, s_{i+1}) + 4 \leq d(x_{\ell-2}, s_j) - d(x_{\ell-2}, s_{i+1}) \quad (1)$$

Let $m \in \{1, 2\}$ be so that $r - m$ is even. Let $v = x_m$. By repeating Equation 1 we get:

$$d(x_{r-2}, s_j) - d(x_{r-2}, s_{i+1}) + 4 \cdot \left(\frac{r}{2} - O(1)\right) \leq d(v, s_j) - d(v, s_{i+1}).$$

Since the distance between any two vertices along the infinite face is at most $\frac{k}{2}$, it follows that: $-\frac{k}{2} + 4 \cdot \left(\frac{r}{2} - O(1)\right) \leq d(v, s_j) - d(v, s_{i+1}) \leq \frac{k}{2}$. Therefore $r \leq \frac{k}{2} + O(1)$. \square

4 A $\Theta(k^2)$ Proof for Halin Graphs

In this section we prove Theorem 4. Namely, we show a tight $\Theta(k^2)$ bound on the number of patterns in a family of graphs that includes Halin graphs. The *Halin graph* family (see [28] for history and properties) is a restricted family of planar graphs. A Halin graph is obtained from an embedded tree \mathcal{T} with no degree-2 vertices by attaching a cycle C to its leaves in their order of appearance according to the embedding. The cycle is then the boundary of the infinite face, and we denote its size by k . We will consider a more general family than Halin graphs. Namely, we allow the tree \mathcal{T} to have degree 2 vertices, and we allow the cycle C to contain vertices that are not in \mathcal{T} . We will refer to such graphs as *S-Halin* graphs. See Figure 13.

In this section, we show that the number of distinct patterns in S-halin graphs is only $O(k^2)$, and that this bound is tight. In contrast, we show that the VC-dimension argument is limited to proving $O(k^3)$ even on such graphs.

Lemma 18. *Let $G = \mathcal{T} \cup C$ be an S-Halin graph, obtained by identifying the leaves of \mathcal{T} with a subset of vertices of a cycle C . If the size of C is k then the number of distinct patterns in G is $O(k^2)$.*

Proof. Notice that the total number of vertices of G with degree larger than 2 is $O(k)$ (and hence, the number of faces in G is also $O(k)$). This is because C is of size k and \mathcal{T} contains only k leaves (and hence $O(k)$ vertices with degree larger than 2). Now consider the dual graph $G^* = (V^*, E^*)$. Since every bisector is a simple cycle in G^* , and since $|V^*| = O(k)$, then the number of arcs in the graphs $G_{\mathcal{B}}$ and $G_{\mathcal{P}}$ is $O(k^2)$. Therefore, by Corollary 12 the number of distinct patterns in G is $O(k^2)$. \square




Figure 13: An S-Halin graph whose matrix of patterns has VC-dimension 3.

We next show that it is not possible to prove Lemma 18 using the VC-dimension argument. Namely, consider the S-Halin graph of Figure 13 and let P be the matrix whose rows are the patterns of the graph (recall that each pattern is in $\{-1, 1\}^{k-1}$).

Proposition 19. *The VC-dimension of P is 3.*

Proof. Consider the following submatrix of P whose rows correspond to the vertices v_0, \dots, v_7 and columns correspond to the edges e_1, e_2, e_3 :

$$\begin{array}{c} \\ v_0 \\ v_1 \\ v_2 \\ v_3 \\ v_4 \\ v_5 \\ v_6 \\ v_7 \end{array} \begin{pmatrix} e_1 & e_2 & e_3 \\ 1 & 1 & 1 \\ 1 & 1 & -1 \\ 1 & -1 & 1 \\ 1 & -1 & -1 \\ -1 & 1 & 1 \\ -1 & 1 & -1 \\ -1 & -1 & 1 \\ -1 & -1 & -1 \end{pmatrix}$$

Since it contains all possible rows, the VC-dimension of P is at least 3. □

It is important to remark that we can generalize the example of Figure 13 to any large enough k , by adding vertices along the infinite face in the part between v_7 and v_0 (clockwise).

So far we have seen that S-Halin graphs have at most $O(k^2)$ distinct patterns (Lemma 18), and that the VC-dimension argument is limited to showing $O(k^3)$ distinct patterns (Proposition 19). To conclude this section, we prove that the $O(k^2)$ bound is tight:

Lemma 20. *There exists an S-Halin graphs with $\Omega(k^2)$ distinct patterns.*

Proof. We assume that k is even and denote $k' = k/2$. We construct the tree \mathcal{T} by taking the union of $k' + 1$ simple paths $P_0, P_1, \dots, P_{k'}$ where every P_i is of length i and all P_i 's originate from

a common vertex $v_{0,0}$. Namely, $P_i = (v_{i,0} - v_{i,1} - \dots - v_{i,i})$, and $v_{i,0} = v_{j,0}$ for every $i \neq j$. We choose the embedding of \mathcal{T} so that in a clockwise tour around $v_{0,0}$, the order of appearance of the paths is $(P_1, P_2, \dots, P_{k'})$.

We define an additional $v_{k',k'}$ -to- $v_{1,1}$ path of length $k' + 1$ denoted $Q = (v_{k',k'} - q_1 - q_2 - \dots - q_{k'} - v_{1,1})$. Let C be the cycle $Q \circ (v_{1,1} - v_{2,2} - \dots - v_{k',k'})$. Let $G = \mathcal{T} \cup C$. See Figure 14. Note that $|C|$ is $2k' = k$.




Figure 14: The S-Halin of Lemma 20 for $k = 8$.

Consider the patterns of G , when we choose the first vertex to be $v_{1,1}$, the second $v_{2,2}$, etc. Let $i \in [k']$, $j \in [i - 1]$. We claim that the pattern of $v_{i,j}$ is:

$$p_{v_{i,j}} = 1^{i-j-1} \circ (-1)^j \circ 1^{k'+j+1-i} \circ (-1)^{k'-j-1} \quad (2)$$

To see why, consider the $v_{i,j}$ -to- $v_{t,t}$ distances for $t \leq i$. Notice that for every $1 \leq t \leq i - j$ we have $d(v_{i,j}, v_{t,t}) = j + t$ and for every $i - j \leq t \leq i$ we have $d(v_{i,j}, v_{t,t}) = 2i - j - t$. In particular, they are equal at $t = i - j$. Therefore, the pattern of all the edges between $v_{1,1}$ and $v_{i,i}$ is $1^{i-j-1} \circ (-1)^j$.

Now consider the $v_{i,j}$ -to- $v_{t,t}$ distances for $i \leq t \leq k'$. Notice that a shortest $v_{i,j}$ -to- $v_{t,t}$ path will never use Q and instead will go through $(v_{i,i} - v_{i+1,i+1} - \dots - v_{t,t})$. Namely, the distance is $d(v_{i,j}, v_{t,t}) = t - j$. Therefore, the pattern of all the edges between $v_{i,i}$ and $v_{k',k'}$ is $1^{k'-i}$.

Finally, consider the $v_{i,j}$ -to- q_t distances for $1 \leq t \leq k'$. For every $1 \leq t \leq j + 1$ we have $d(v_{i,j}, q_t) = k' + t - j$ and for every $j + 1 \leq t \leq k'$ we have $d(v_{i,j}, q_t) = k' + j + 2 - t$. In particular they are equal at $t = j + 1$. Therefore the pattern of all the edges between $v_{k',k'}$ and q'_k is $1^{j+1} \circ (-1)^{k'-j-1}$. Overall, we get that $p_{v_{i,j}}$ is as in Equation (2).

Notice that $p_{v_{i,j}}$ is unique for every different $i \in [k']$, $j \in [i - 1]$. Since the number of such vertices is $\Omega(k'^2) = \Omega(k^2)$, there are $\Omega(k^2)$ distinct patterns in G . \square

5 Conclusions

In this work we developed a technique for analyzing the structure and number of distinct patterns in undirected planar graphs. This technique leads to an improved $\tilde{O}(p_{\#} + k + |T|)$ space compression of the Okamura-Seymour metric, and to an optimal $\tilde{O}(k + |T|)$ compression in the special case where the vertices of T induce a connected component in G . Moreover, the technique leads to an alternative proof of the $p_{\#} = O(k^3)$ upper bound on the number of different patterns.

We have shown that for the family of Halin graphs, the original proof technique using VC-dimension is not tight, and that our approach easily proves the tight bound in this case. Going

back to planar graphs, we were unable to come up with constructions of families of planar graphs that have $p_{\#} = \omega(k^2)$ patterns. We therefore make the following conjecture.

Conjecture. *The number of distinct patterns over all vertices of a planar graph is $O(k^2)$.*

We hope that tools we have developed in this work will be useful in proving this conjecture.

References

- [1] Amir Abboud, Pawel Gawrychowski, Shay Mozes, and Oren Weimann. Near-optimal compression for the planar graph metric. In *29th SODA*, pages 530–549, 2018.
- [2] Amitabh Basu and Anupam Gupta. Steiner point removal in graph metrics, 2008. Unpublished Manuscript, available from <http://www.math.ucdavis.edu/~abasu/papers/SPR.pdf>.
- [3] Jon Louis Bentley. Solutions to klee’s rectangle problems. *Unpublished manuscript*, pages 282–300, 1977.
- [4] Guy E. Blelloch and Arash Farzan. Succinct representations of separable graphs. In *21st CPM*, pages 138–150, 2010.
- [5] Hubert T.-H. Chan, Donglin Xia, Goran Konjevod, and Andréa W. Richa. A tight lower bound for the steiner point removal problem on trees. In *9th APPROX*, pages 70–81, 2006.
- [6] Hsien-Chih Chang, Pawel Gawrychowski, Shay Mozes, and Oren Weimann. Near-optimal distance emulator for planar graphs. In *26th ESA*, pages 16:1–16:17, 2018.
- [7] Hsien-Chih Chang and Tim Ophelders. Planar emulators for monge matrices. In *32nd CCCG*, pages 141–147, 2020.
- [8] Yun Kuen Cheung. Steiner point removal - distant terminals don’t (really) bother. In *29th SODA*, pages 1353–1360, 2018.
- [9] Yun Kuen Cheung, Gramoz Goranci, and Monika Henzinger. Graph minors for preserving terminal distances approximately - lower and upper bounds. In *43rd ICALP*, pages 131:1–131:14, 2016.
- [10] Yi-Ting Chiang, Ching-Chi Lin, and Hsueh-I Lu. Orderly spanning trees with applications to graph encoding and graph drawing. In *22nd SODA*, pages 506–515, 2001.
- [11] James R Driscoll, Neil Sarnak, Daniel D Sleator, and Robert E Tarjan. Making data structures persistent. *Journal of computer and system sciences*, 38(1):86–124, 1989.
- [12] David Eisenstat and Philip N Klein. Linear-time algorithms for max flow and multiple-source shortest paths in unit-weight planar graphs. In *45th STOC*, pages 735–744, 2013.
- [13] Arnold Filtser. Steiner point removal with distortion $O(\log k)$. In *29th SODA*, pages 1361–1373, 2018.
- [14] Arnold Filtser, Robert Krauthgamer, and Ohad Trabelsi. Relaxed voronoi: A simple framework for terminal-clustering problems. In *2nd SOSA*, pages 10:1–10:14, 2019.

- [15] Viktor Fredslund-Hansen, Shay Mozes, and Christian Wulff-Nilsen. Truly subquadratic exact distance oracles with constant query time for planar graphs. In *32nd ISAAC*, 2021. To appear.
- [16] Cyril Gavoille, David Peleg, Stéphane Pérennes, and Ran Raz. Distance labeling in graphs. *Journal of Algorithms*, 53(1):85–112, 2004.
- [17] Gramoz Goranci, Monika Henzinger, and Pan Peng. Improved guarantees for vertex sparsification in planar graphs. In *25th ESA*, pages 44:1–44:14, 2017.
- [18] Anupam Gupta. Steiner points in tree metrics don’t (really) help. In *12th SODA*, pages 220–227, 2001.
- [19] Lior Kamma, Robert Krauthgamer, and Huy L. Nguyen. Cutting corners cheaply, or how to remove steiner points. *SIAM Journal of Computing*, 44(4):975–995, 2015.
- [20] Richard M. Karp and Michael O. Rabin. Efficient randomized pattern-matching algorithms. *IBM Journal of Research and Development*, 31(2):249–260, 1987.
- [21] Ken-ichi Kawarabayashi, Philip N Klein, and Christian Sommer. Linear-space approximate distance oracles for planar, bounded-genus and minor-free graphs. In *38th ICALP*, pages 135–146, 2011.
- [22] Philip N Klein. Multiple-source shortest paths in planar graphs. In *16th SODA*, volume 5, pages 146–155, 2005.
- [23] Robert Krauthgamer, Huy L Nguyen, and Tamar Zondiner. Preserving terminal distances using minors. *SIAM Journal on Discrete Mathematics*, 28(1):127–141, 2014.
- [24] Robert Krauthgamer and Tamar Zondiner. Preserving terminal distances using minors. In *39th ICALP*, pages 594–605, 2012.
- [25] Jason Li and Merav Parter. Planar diameter via metric compression. In *51st STOC*, pages 152–163, 2019.
- [26] J. Ian Munro and Venkatesh Raman. Succinct representation of balanced parentheses, static trees and planar graphs. In *38th FOCS*, pages 118–126, 1997.
- [27] Norbert Sauer. On the density of families of sets. *Journal of Combinatorial Theory, Series A*, 13(1):145–147, 1972.
- [28] Maciej M Sysło and Andrzej Proskurowski. On halin graphs. In *Graph Theory*, pages 248–256. Springer, 1983.
- [29] Mikkel Thorup. Compact oracles for reachability and approximate distances in planar digraphs. *Journal of the ACM*, 51(6):993–1024, 2004.
- [30] György Turán. On the succinct representation of graphs. *Discrete Applied Mathematics*, 8(3):289–294, 1984.