Transforming Spanning Trees: A Lower Bound

Kevin Buchin[*] Andreas Razen[†] Takeaki Uno[‡] Uli Wagner[†]

Abstract

For a planar point set we consider the graph of crossing-free straight-line spanning trees where two spanning trees are adjacent in the graph if their union is crossing-free. An upper bound on the diameter of this graph implies an upper bound on the diameter of the flip graph of pseudo-triangulations of the underlying point set.

We prove a lower bound of $\Omega\big(\log(n)/\log(\log(n))\big)$ for the diameter of the graph of spanning trees on a planar set of n points. This nearly matches the known upper bound of $O(\log(n))$. If we measure the diameter in terms of the number of convex layers k of the point set, our lower bound construction is tight, i.e., the diameter is in $\Omega(\log(k))$ which matches the known upper bound of $O(\log(k))$. So far only constant lower bounds were known.

1 Introduction

Given a set S of n points in the plane let T_S denote the set of all crossing-free straight-line spanning trees of S. A straight-line embedded graph is *crossing-free* if every pair of its edges does not share any point other than common endpoints. We call two crossing-free spanning trees T_1 and T_2 of S *compatible* if their union, i.e. the graph on S with edge set $E(T_1) \cup E(T_2)$, is crossing-free.

Aichholzer, Aurenhammer, and Hurtado [2] investigate how fast two spanning trees can be transformed into each other by a sequence of spanning trees with any two consecutive trees being compatible. They prove that the maximum length of a sequence needed is in $O(\log(n))$.

Let $\mathcal{T}_S$ denote the graph with T_S as vertex set and edges between compatible spanning trees. The maximum length of a sequence needed to transform two spanning trees corresponds to the diameter of this graph. Aichholzer, Aurenhammer, Huemer, and Krasser [1] refine the above bound on the diameter of $\mathcal{T}_S$ to a bound of $O(\log(k))$, where k denotes the number of convex layers of S. The *convex layers* of a point set S are defined inductively: the first convex layer U_1 consists of the boundary points of the convex hull of S and, for $i > 1$, the i-th convex layer U_i is defined as the set of boundary points of the convex hull of $S \setminus \bigcup_{j<i} U_j$. The number k of convex layers of a point set S is the minimum i such that $U_{i+1} = \emptyset$.

Aichholzer et al. [1] also prove that an upper bound of d on the diameter of $\mathcal{T}_S$ yields an upper bound of $O(nd)$ on the diameter of the flip graph of pseudo-triangulations of S. They conjecture that the diameter of $\mathcal{T}_S$ is sublogarithmic. So far no example was known where the diameter is not constant.

We give a sublogarithmic but considerably tighter lower bound: we complement the $O(\log(n))$ upper bound with a lower bound of $\Omega\big(\log(n)/\log(\log(n))\big)$. We do this constructively by providing point sets of increasing size, and on each point set we specify two spanning trees achieving this bound. For these examples the bound in the number of convex layers is tight, i.e., the distance between the two trees is in $\Omega(\log(k))$, where k is the number of convex layers.

2 The Lower Bound

In this section we construct point sets in the plane and consider pairs of spanning trees which need a large number of transformation steps to transform one tree into the other.

We will first develop a general scheme to construct such trees. Based on this we present two recursive constructions using this scheme in different ways. The first yields a lower bound of $\Omega(\sqrt{\log(n)})$ on the number of transformations needed, where n is the size of the point set. The second gives a lower bound of $\Omega\big(\log(n)/\log(\log(n))\big)$. Both constructions use point sets with more than two points on a line, i.e., the points are not in general position. However, they can easily be changed to do so by applying a small perturbation, without losing any of the relevant properties of the construction.

The basic concept of the constructions is that by placing the topmost vertex of the point set very far away from the others, we consider a first tree with only near vertical edges and a second tree with many near horizontal edges crossing the vertical edges of the first tree. Furthermore, there are dependencies between the horizontal edges such that, when transforming one tree into the other, a vertex that connects to the rest of the tree by a horizontal edge may connect to the topmost vertex by a vertical edge only

---

[*] Institute of Computer Science, Freie Universität Berlin, `buchin@inf.fu-berlin.de`

[†] Institute of Theoretical Computer Science, ETH Zurich, `{arazen, wagneru}@inf.ethz.ch`

[‡] National Institute of Informatics, Tokyo, Japan, `uno@nii.jp`

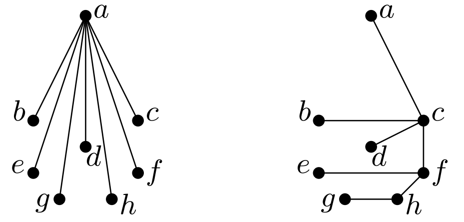

(a) T_1 (b) T_2

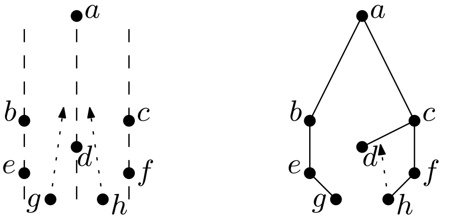

(c) Two vertical strips defined by $b, c, \dots, f$ (d) One vertical strip is blocked

Figure 1: The trees T_1 and T_2 have distance 3 in $\mathcal{T}_S$.

if certain other horizontal edges are no longer in the tree.

We illustrate this by the example in Figure 1 with $S = \{a, b, \dots, h\}$ being the underlying point set. The first tree T_1 (Figure 1(a)) has near vertical edges, the second tree T_2 (Figure 1(b)) has mostly near horizontal edges. The points b, c, d, e, and f subdivide the space in which the points of S lie into two vertical strips. In each such strip there is one point at the bottom (g and h) which needs to connect to the topmost point a through the corresponding strip (Figure 1(c)). At the beginning the edges $\{b, c\}$ and $\{e, f\}$ block both strips completely, i.e., the bottommost points g and h cannot connect to a in any neighbor of T_2 in $\mathcal{T}_S$. Furthermore, whatever the first transformation is, thereafter the point d will have an edge to at least one of b, c, e, or f (as in the example of the tree in Figure 1(d)). Thus, after one transformation the edge $\{a, g\}$ or $\{a, h\}$ still crosses an edge of the current tree and cannot be present in the next transformation. In total, three transformations are necessary and also suffice to transform T_2 to T_1, and the diameter of $\mathcal{T}_S$ is at least 3.

2.1 Blocking Vertical Strips

Before turning to the construction of a point set, we further develop the concept of blocking vertical strips. A *vertical strip* R is a subset of $\mathbb{R}^2$ such that there exist $a, b \in \mathbb{R}$ with

$$R = \{(x, y) \in \mathbb{R}^2 | a \leq x \leq b\} =: [a, b] \times \mathbb{R};$$

the *width* of the vertical strip R is $b - a$. An edge *blocks a vertical strip* if the end points of the edge lie on different sides (possibly on the border) of the strip. For instance, in Figure 1(b) (assuming a proper coordinate system) the edges $\{b, c\}$ and $\{e, f\}$ of the tree T_2 both block the vertical strip $[0, 1] \times \mathbb{R}$ and the edge $\{d, c\}$ blocks the vertical strip $[1/2, 1] \times \mathbb{R}$.

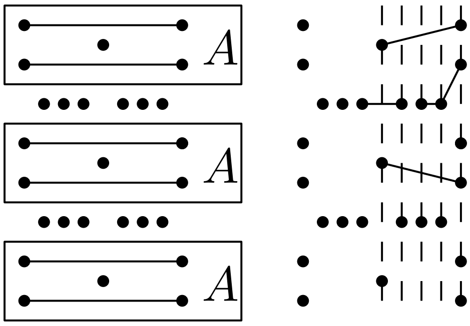

(a) Stacking three 1-of-2-blockers (b) After 1 transformation

Figure 2: A 3-of-8-blocker after 2 steps.

A point set A together with a set E of straight-line edges on A *blocks a vertical strip of width* $w > 0$ *after* k *steps*, if for any point set S containing A (and no further point in the convex hull of A) the following holds: if a spanning tree $T \in T_S$ contains the edges E then in any spanning tree in the k–neighborhood of T in $\mathcal{T}_S$ some vertical strip of width at least w is blocked (not necessarily by an edge in E). For instance, in T_2 in Figure 1(b) the points b, c, d, e, and f together with the edges $\{b, c\}$ and $\{e, f\}$ block a vertical strip of width 1/2 after 1 step: either the strip $[0, 1/2] \times \mathbb{R}$ is blocked by the edge $\{b, d\}$ or $\{e, d\}$, or the strip $[1/2, 1] \times \mathbb{R}$ is blocked by $\{d, c\}$ or $\{d, f\}$.

Note that this concept now implies the following: assume that we have a point set S with the topmost point $p_0 \in S$ placed very far away from the rest, and $A \subset S$ with edges E on A blocks some vertical strip R after k steps. Let $T_1 \in T_S$ be the tree where p_0 connects to every other point by a near vertical edge and let $T_2 \in T_S$ contain the edges E. If there is a point in $S \cap R$ lying strictly below the edge responsible for blocking R after k steps then T_1 cannot be in the $(k+1)$–neighborhood of T_2 in $\mathcal{T}_S$. Thus, the diameter of $\mathcal{T}_S$ is at least $k + 2$.

The point sets we are about to construct reside in the strip $[0, 1] \times \mathbb{R}$, and therein we consider specific vertical strips that might be blocked. We call a point set A together with a set of edges E an *l-of-m-blocker after k steps* if A blocks at least l of the vertical strips $[(i-1)/m, i/m] \times \mathbb{R}$ (for $i = 1, \dots, m$) after k steps, not necessarily the same strips for different trees containing E in their respective k–neighborhood in $\mathcal{T}_S$. In the example of T_2 in Figure 1(b) the points b, c, d, e, f together with the edges $\{b, c\}, \{e, f\}$ are a 1-of-2-blocker after 1 step. We call l/m the *density* of the blocker.

Given a blocker we can construct further blockers with a larger number of steps by stacking the blocker and spreading in further points. Consider for instance the construction in Figure 2(a). It contains three copies of the 1-of-2-blocker after 1 step, A, together with the corresponding horizontal edges, and between two adjacent copies of A additional points subdivide each (of two) strips into four smaller strips resulting in a total of eight vertical strips. After 1 step each copy of A blocks one vertical strip of width $1/2$. Since there are three copies of A by the pigeon-hole principle one strip is blocked twice (in the example, Figure 2(b), the right vertical strip). No matter how the points in-between these blocking edges are connected to the rest of the tree at least three of the four corresponding vertical strips of width $1/8$ are blocked, and this can only change after the edges blocking the strip of width $1/2$ are removed. This is the case at the earliest after 2 steps, thus the construction is a 3-of-8-blocker after 2 steps, and for a point set S containing this blocker the diameter of $\mathcal{T}_S$ is at least 4.

This construction is generalized in the following.

Lemma 1 *Let A be an l-of-m-blocker after k steps of density $l/m > 1/u$ for some $u \in \mathbb{N}$. Stacking u copies of A on top of each other with additional points (equidistantly) subdividing each of the m vertical strips into m' vertical strips between each pair of adjacent copies yields an $(m'-1)$-of-$(m \cdot m')$-blocker after $k+1$ steps.*

Proof. After k steps the u copies of A block within the m vertical strips $l \cdot u > m$ times, thus at least one of the m strips is blocked twice. The points in this vertical strip blocked from above and below subdivide this strip into m' smaller strips, hence in order to connect these points to the rest at least $m'-1$ of the small strips are blocked. This changes at the earliest after $k+1$ steps, thus the construction is an $(m'-1)$-of-$(m \cdot m')$-blocker after $k+1$ steps. ☐

2.2 Construction 1

We construct a point set S depending on an integer variable d together with two trees $T_1, T_2 \in T_S$ such that at least d steps are needed to transform one of the trees into the other, and the size of S is in $O(2^{d^2})$, i.e., $d \in \Omega(\sqrt{\log(n)})$, where $n = |S|$.

All points of S lie in the infinite strip $[0,1] \times \mathbb{R}$. A special point p_0 has a larger y-coordinate than all other points, and will be chosen such that the slope of any line through p_0 and any other point in S is larger than the slopes of all non-vertical lines through two points in $S \setminus p_0$.

Let L_0 be defined as $L_0 := \{(0,0),(1,0)\}$ and L_k for $k \in \mathbb{N}, k \geq 1$ as

$$L_k := \left\{ \left(\frac{2i-1}{2^k}, 0\right) \middle| i = 1, \ldots, 2^{k-1} \right\}.$$

Thus, $\bigcup_{0 \leq k' \leq k} L_{k'}$ subdivides the line segment from $(0,0)$ to $(1,0)$ into 2^k equal parts by 2^k+1 points. The set L_{k+1} places one point in the center of each of these parts.

We define point sets A_k, $k \in \mathbb{N}$, inductively. Let

$$A_1 := L_0 \ \cup \ L_1 \oplus 1 \ \cup \ L_0 \oplus 2,$$

where $P \oplus t := \{(x, y+t) | (x,y) \in P\}$ is a vertical shift of the point set $P \subset \mathbb{R}^2$ by $t \in \mathbb{N}$. Note that A_1 corresponds to the point set A from Figure 2(a).

For $k \in \mathbb{N}$, let A_{k+1} be defined by stacking 2^k+1 copies of A_k with a copy of L_{k+1} between each pair of adjacent copies of A_k. Formally,

$$\begin{aligned} A_{k+1} \quad := \quad & \bigcup_{i=0}^{2^k} A_k \oplus i \cdot (h_k + 1) \\ & \cup \bigcup_{i=0}^{2^k-1} L_{k+1} \oplus \big(i \cdot (h_k+1) + h_k\big), \end{aligned}$$

where $h_k := 2 \cdot \prod_{i=0}^{k-1} (2^i + 1) - 1$.

It follows directly from Lemma 1 that the point set A_k together with edges between every pair of points with coordinates $(0,y),(1,y)$, for some $y \in \mathbb{N}$, is a 1-of-2^k-blocker after k steps.

Given $d \in \mathbb{N}$ define $S := L_{d+1} \cup A_d \oplus 1 \cup \{p_0\}$ with p_0 chosen as described above. Let T_1 be the star connecting p_0 to every other point by an edge. Let T_2 be a tree on S obtained by taking all (exactly) horizontal edges blocking the complete vertical strip $[0,1] \times \mathbb{R}$ and adding further edges such that T_2 is a crossing-free straight-line spanning tree. We already know that A_d together with the corresponding horizontal edges is a 1-of-2^d-blocker after d steps. Thus, when transforming T_2 into T_1 there will be one of the points in L_{d+1} blocked away from p_0 after d steps. Therefore, at least $d+2$ transformations are needed.

The cardinality s_d of A_d is given by $s_1 = 5$ and the recursion $s_{k+1} = (2^k+1)s_k + 2^k \cdot 2^k$. Thus, we have $s_{k+1} \leq 2^{2k+1} s_k$ and by induction $s_d \leq 5 \cdot 2^{d^2}$. The size of S is $2^d + s_d + 1$, hence $d \in \Omega(\sqrt{\log(|S|)})$.

Next we consider the number of convex layers. The first layer of S consists of the topmost point, the points of the bottom row, the points in the left most and the right most column of points. With each additional convex layer two more rows and two more columns are considered until only one row or one column is left. If m_1 is the number of different x-coordinates used and m_2 the number of different y-coordinates used in the construction then we can bound the number of convex layers from above by

$$1 + \frac{1}{2} \min(m_1, m_2).$$

The number of different x-coordinates in S is bounded by 2^d, thus d is logarithmic in the number of convex layers.

Theorem 2 *There is a point set S in the plane for which the diameter of $\mathcal{T}_S$ is in $\Omega(\log(k))$, where k is the number of convex layers of S.*

2.3 Construction 2

The point set A_k from Construction 1 suffered from an exponential growth in both, the number of copies of A_{k-1} and the number of points in L_k placed in-between. Note that the recursive construction we present in the following will only require a constant number of copies of previously constructed point sets.

We construct a point set $S \subset [0,1] \times \mathbb{R}$ depending on an integer variable $d > 1$ together with two trees $T_1, T_2 \in \mathcal{T}_S$ such that $d \in \Omega\big(\log(n)/\log(\log(n))\big)$, where n is the size of S, and the distance of the trees in $\mathcal{T}_S$ is at least $\lfloor d/2 \rfloor$.

Again, a special point p_0 is included in S with a far larger y-coordinate than any other point in S.

However, contrary to the first construction where the density of the blockers dropped by a factor of $1/2$ in every step, we will now keep the density above $1/2$ as long as possible by spreading in more points. For this purpose let $L_0 := \{(0,0),(1,0)\}$, and for $k \geq 1$ define

$$L_k := \left\{ (i/d^{k-1} + j/d^k, 0) \,\middle|\, \begin{array}{c} i = 0, \dots, d^{k-1}-1 \\ j = 1, \dots, d-1 \end{array} \right\},$$

i.e., $\bigcup_{0 \leq k' \leq k} L_{k'}$ subdivides the line segment from $(0,0)$ to $(1,0)$ into d^k equal parts by $d^k + 1$ points.

We define the point sets A_k inductively. Let

$$A_1 := L_0 \ \cup\ L_1 \oplus 1 \ \cup\ L_0 \oplus 2,$$

and for $k \in \{1, \dots, \lfloor d/2 \rfloor - 1\}$ and $h_k := 4 \cdot 3^{k-1} - 1$,

$$\begin{aligned} A_{k+1} \quad := \quad & A_k \ \cup\ L_{k+1} \oplus h_k \ \cup\ A_k \oplus (h_k + 1) \\ & \cup\ L_{k+1} \oplus (2h_k + 1) \ \cup\ A_k \oplus (2h_k + 2). \end{aligned}$$

Note that here A_{k+1} only uses three copies of the previously constructed A_k.

The point set A_1 and the horizontal edges between points with coordinates $(0,y)$ and $(1,y)$, for some $y \in \mathbb{N}$, form a $(d-1)$-of-d-blocker after 1 step. Applying Lemma 1 at this time gives that A_k together with the corresponding edges is a $(d-1)$-of-d^k-blocker after k steps. However, taking a closer look we can prove something stronger: recall that for the blocker A_1 at most one vertical strip of width $1/d$ is not blocked after 1 step. Placing three copies of A_1 on top of each other implies that after 1 step there cannot be more than one vertical strip of width $1/d$ that is not blocked at least twice. Hence, each of the $d-1$ vertical strips of width $1/d$ that are blocked twice, together with the points from L_2 in-between, behave like a (horizontally) scaled blocker A_1.

See for instance Figure 3 with the corresponding construction for $d = 4$. In Figure 3(b) only blocking edges (of the scaled blockers) are drawn as solid lines.

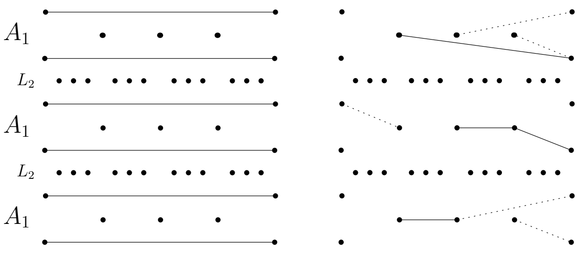

(a) The blocker A_2 and the horizontal edges, for $d = 4$

(b) Blocking edges (drawn as solid lines) after 1 step

Figure 3: A_2 is a 9-of-16-blocker after 2 steps.

Therefore, A_2 together with the horizontal edges is a $(d-1)^2$-of-d^2-blocker after 2 steps (since $d \geq 2$).

Inductively we find that A_k with the corresponding edges is a $(d-1)^k$-of-d^k-blocker after k steps as long as the density $(d-1)^{k-1}/d^{k-1}$ of the blocker A_{k-1} is at least $1/2$ such that the three copies suffice to guarantee the existence of some blocked vertical strip.

As $d \geq 2$, this holds for $k = \lfloor d/2 \rfloor$. Thus, $A_{\lfloor d/2 \rfloor}$ is a blocker with density at least $1/2$ after $\lfloor d/2 \rfloor$ rounds. With $S := L_{\lfloor d/2 \rfloor + 1} \cup A_{\lfloor d/2 \rfloor} \oplus 1 \cup \{p_0\}$ and T_1 and T_2 defined as in the first construction, the distance of the two trees in $\mathcal{T}_S$ is $\lfloor d/2 \rfloor + 2$. The size s_1 of A_1 is $d + 3$ and the size s_{k+1} of A_{k+1} can be bounded by the recursion $s_{k+1} < 3 \cdot s_k + 2 \cdot d^{k+1}$. For $d \geq 3$ we get by induction $s_k < 2kd^k$. This yields $|S| < d^{\lfloor d/2 \rfloor + 1} + 2\lfloor d/2 \rfloor d^{\lfloor d/2 \rfloor} + 1$ and hence $d \in \Omega\big(\log(|S|)/\log(\log(|S|))\big)$.

To express the diameter of the set S in terms of the number of convex layers we use the same argument as in Construction 1 but now count the rows instead of the columns. The number of rows is of order $3^{\lfloor d/2 \rfloor}$, thus the diameter is again logarithmic in the number of convex layers.

Theorem 3 *There exists a set S of n points in the plane for which the diameter of the graph $\mathcal{T}_S$ is in $\Omega\big(\log(n)/\log(\log(n))\big)$.*

We have the feeling that the $1/\log(\log(n))$ factor in the lower bound from Theorem 3 is more likely to be an artifact of our construction than the truth about the diameter of $\mathcal{T}_S$ which we think should be in $\Theta(\log(n))$, for a suited S with $n = |S|$.

References

[1] O. Aichholzer, F. Aurenhammer, C. Huemer, and H. Krasser. Transforming spanning trees and pseudo-triangulations. *Inf. Process. Lett. 97*, 1 (2006), 19–22.

[2] O. Aichholzer, F. Aurenhammer, and F. Hurtado. Sequences of spanning trees and a fixed tree theorem. *Comput. Geom. Theory Appl. 21*, 1 (2002), 3–20.