

Improved Bounds for the Union of Locally Fat Objects in the Plane*

Boris Aronov[†] Mark de Berg[‡] Esther Ezra[§] Micha Sharir[¶]

October 24, 2013

Abstract

We show that, for any $\gamma > 0$, the combinatorial complexity of the union of n locally γ -fat objects of constant complexity in the plane is $\frac{n}{\gamma^4} 2^{O(\log^* n)}$. For the special case of γ -fat triangles, the bound improves to $O(n \log^* n + \frac{n}{\gamma} \log^2 \frac{1}{\gamma})$.

*Work on this paper by Boris Aronov and Micha Sharir has been supported by a joint grant No. 2006/194 from the U.S.-Israel Binational Science Foundation. Work by Boris Aronov has also been supported by NSA MSP Grant H98230-10-1-0210 and NSF Grants CCF-08-30691 and CCF-11-17336. Work by Micha Sharir has also been supported by NSF Grant CCF-08-30272, by Grants 338/09 and 892/13 from the Israel Science Foundation, by Grant 2012/229 from the US-Israel Binational Science Foundation, by the Israeli Centers for Research Excellence (I-CORE) program (center no. 4/11), and by the Hermann Minkowski–MINERVA Center for Geometry at Tel Aviv University. Mark de Berg was supported by the Netherlands' Organisation for Scientific Research (NWO) under project no. 639.023.301. Work on this paper by Esther Ezra has been supported by NSF under grant CCF-12-16689. A preliminary version of this paper has appeared in *Proc. Sympos. Discrete Algorithms*, 2011, pp. 1778–1785 [21].

[†]Department of Computer Science and Engineering, Polytechnic Institute of NYU, Brooklyn, NY 11201-3840, USA; aronov@poly.edu.

[‡]Department of Computing Science, TU Eindhoven, P.O. Box 513, 5600 MB Eindhoven, the Netherlands; mdberg@win.tue.nl.

[§]Courant Institute of Mathematical Sciences, New York University, New York, NY 10012, USA; esther@courant.nyu.edu.

[¶]School of Computer Science, Tel Aviv University, Tel Aviv 69978 Israel and Courant Institute of Mathematical Sciences, New York University, New York, NY 10012, USA; michas@post.tau.ac.il.

1 Introduction

In this paper we obtain sharper upper bounds on the complexity of the union of n locally γ -fat objects of constant complexity in the plane, and of n γ -fat triangles in the plane; see below for the definitions of these classes of objects and for the precise statements of our bounds.

Background. Consider a family \mathcal{F} of well-behaved and simply-shaped geometric objects in the plane; we will formally refer to them as objects of *constant complexity*, and give a precise definition of this notion below. For now assume \mathcal{F} to consist of Jordan regions with interiors so that every pair of boundaries intersect in at most some fixed constant number of points. We denote the union of \mathcal{F} by $\mathcal{U}(\mathcal{F})$. The (*combinatorial*) *complexity* of $\mathcal{U}(\mathcal{F})$, which we denote by $|\mathcal{U}(\mathcal{F})|$, is defined as the total number of vertices of the union boundary, which can either be vertices of original objects or intersections between pairs of object boundaries. In this paper we are interested in the maximum complexity of the union when the objects of \mathcal{F} are *fat*, as defined below.

There are many algorithms and data structures whose performance depends on the union complexity of some set of geometric objects. Hence, the problem of bounding the union complexity of certain types of geometric objects has received considerable attention over the past 25 years; see the recent survey of Agarwal *et al.* [1] for an overview and many references. In general, the union complexity of a set \mathcal{F} of n planar objects of constant complexity can be $\Theta(n^2)$ —this is achieved, for example, when \mathcal{F} consists of n long and thin triangles arranged in a grid-like pattern—but for several special cases the union complexity is much smaller. One such case is when \mathcal{F} is a set of *pseudodisks*, that is, when the objects in \mathcal{F} are simply connected and the boundaries of any pair of objects in \mathcal{F} intersect in at most two points. Kedem *et al.* [28] showed that in this case the maximum union complexity is only linear. Another case is when \mathcal{F} is a set of so-called fat objects. We now review in more detail the known results for fat objects in the plane; for results on fat objects in \mathbb{R}^3 , see Ezra and Sharir [22] and references therein; for a fairly up-to-date survey of the general subject, see the aforementioned survey [1] and references therein.

One of the earliest results on the union complexity of fat objects was for planar γ -fat wedges, namely, unbounded wedges, each with an opening angle at least γ , for some fixed constant $\gamma > 0$. Alt *et al.* [2] and Efrat *et al.* [18] showed that the complexity of the union of n such wedges is $O(n)$.¹ Later, Matoušek *et al.* [30, 31] studied the more difficult case of γ -fat triangles, namely, triangles all of whose angles are at least γ . They proved that the union of n γ -fat triangles in the plane has only $O(n)$ holes—where a *hole* is a connected component of the complement of the union—and that the union complexity is $O(n \log \log n)$. Van Kreveld [29] showed that the bound extends to certain types of fat polygons (for a somewhat different notion of “fatness”). The best known lower bound on the maximum union complexity of fat triangles is $\Omega(n\alpha(n))$, where $\alpha(n)$ denotes the functional inverse of the Ackermann function [34]. For almost twenty years no improvements were obtained on the union complexity of fat triangles (except for an improvement of the dependence on γ in the constant of proportionality [32], which is now $O(\frac{1}{\gamma} \log \frac{1}{\gamma})$). In the preliminary version of this paper, Ezra *et al.* [21] have managed to improve the bound in [30, 31] to $O(n2^{\alpha(n)} \log^* n)$. The factor $2^{\alpha(n)}$ was replaced by $\alpha(n)$ by Pettie [33]; we remove this factor entirely in this version of the paper—see the discussion below.

In the meantime there has been a lot of work on the union complexity of a set \mathcal{F} of n curved objects of constant complexity and controlled shape. The first result was obtained by Efrat

¹In bounds like this, the constant of proportionality generally depends on γ .

and Sharir [19], who studied *convex* curved fat objects. They defined a convex planar object o to be γ -fat if there exist two concentric disks D_{in} and D_{out} such that $D_{\text{in}} \subseteq o \subseteq D_{\text{out}}$ and $\text{radius}(D_{\text{out}})/\text{radius}(D_{\text{in}}) \leq \gamma$, where $\gamma > 1$ is a fixed constant. Efrat and Sharir proved that the union complexity of any set \mathcal{F} of n such objects, each of constant complexity, is $O(n^{1+\varepsilon})$, for any fixed $\varepsilon > 0$.² Efrat and Katz [17] obtained a better bound, namely $O(\lambda_{s+2}(n) \log n)$, for so-called κ -curved objects (an object o is κ -curved if for every point p of its boundary there exists a disk D of diameter $\kappa \cdot \text{diam}(o)$ passing through p and contained in o); here s is the maximum number of times any two object boundaries intersect, and $\lambda_t(q)$ is the maximum length of Davenport-Schinzel sequences of order t on q symbols; it is near-linear in q for any fixed t [34]. Unfortunately, the class of κ -curved objects is rather restricted, since it does not allow convex vertices; in particular, fat triangles are not κ -curved. The first general result for fat, curved objects—where the objects can be non-convex and can have convex vertices—was obtained by Efrat [16]. He proved that the union complexity of n so-called (α, β) -covered objects (an object o is (α, β) -covered if, for every point p of its boundary, there exists an α -fat triangle Δ with p as a vertex, contained in o , and having edges of length at least $\beta \cdot \text{diam}(o)$) is $O(\lambda_{s+2}(n) \log^2 n \log \log n)$, with s defined as above. Recently this has been further generalized and improved by De Berg [6, 7], who showed that the union complexity of n so-called locally γ -fat objects is $O(\lambda_{s+2}(n) \log n)$.

The class of locally γ -fat objects is the most general class of fat objects in the plane for which near-linear bounds on the union complexity are known; this is the class that we consider in this paper. It is defined as follows. Let γ be a fixed constant parameter, with $0 < \gamma < 1$. A planar object o is called *locally γ -fat* if, for any disk D whose center lies in o and that does not fully contain o in its interior, we have $\text{area}(D \cap o) \geq \gamma \cdot \text{area}(D)$, where $D \cap o$ is the connected component of $D \cap o$ that contains the center of D .

Our results. We prove that the maximum union complexity of n locally γ -fat objects of constant complexity (see below for precise definition) is $\frac{n}{\gamma^4} 2^{O(\log^* n)}$, where the constant of proportionality in the exponent $O(\log^* n)$ depends on the description complexity of the given objects, but not on γ . This considerably improves the previously best known bound of $O(\lambda_{s+2}(n) \log n)$, mentioned above, by De Berg [7]. Our analysis can be fine-tuned for the case of γ -fat triangles, exploiting better machinery for handling unions of polygonal objects (the so-called Combination Lemma—see below), and yields the improved bound of $O(n \log^* n + \frac{n}{\gamma} \log^2 \frac{1}{\gamma})$, (see Theorem 4.10). This is a strengthening of a very recent $O(n\alpha(n) \log^* n + \frac{n}{\gamma} \log^2 \frac{1}{\gamma})$ upper bound result of Pettie [33],³ which in turn was an improvement on the $O(n2^{\alpha(n)} \log^* n + \frac{n}{\gamma} \log^2 \frac{1}{\gamma})$ bound from the preliminary version of this paper [21].

2 Preliminaries

We begin by restating and expanding some of the definitions made in the introduction. For an angle γ , with $0 < \gamma \leq \pi/3$, a γ -fat triangle is a triangle all of whose angles are at least γ . An *object* is a simply-connected compact set with interior in the plane. For a fixed constant γ , with $0 < \gamma \leq 1/4$, a planar object o is called *locally γ -fat* if, for any disk D whose center lies in o and that does not fully contain o in its interior, $\text{area}(D \cap o) \geq \gamma \cdot \text{area}(D)$, where $D \cap o$ is the connected component of $D \cap o$ containing the center of D ; see Figure 1. To avoid introducing too many symbols, we shall

²In bounds of this form, the constant of proportionality depends on ε , and generally tends to infinity as ε decreases to 0.

³The dependence on γ was not stated explicitly there, but can be deduced from the analysis.

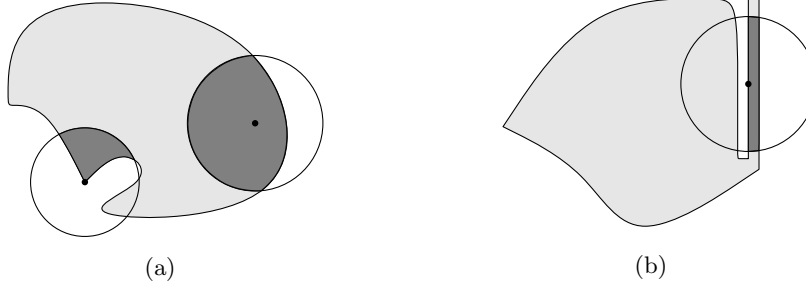


Figure 1: A locally fat object (a), and a fat object that is not locally fat (b).

use the same symbol γ for the two distinct fatness parameters in the above definitions. This is not a significant abuse of notation, since a γ -fat triangle is locally γ' -fat for $\gamma' = \tan \gamma / (4\pi)$.

The definition of local fatness is similar to the fatness definition of Van der Stappen [35], except that his definition uses $D \cap o$ instead of $D \sqcap o$. For convex objects the definitions are thus equivalent, while for non-convex objects local fatness is a stronger condition. This is needed, however, to guarantee small union complexity, as the union complexity of n constant-complexity non-convex objects that are fat under Van der Stappen's definition can be $\Omega(n^2)$ [6].

An object is of *constant (algebraic description) complexity* if it can be described by a Boolean formula constructed from at most c algebraic inequalities in x and y , of total degree at most c , for some constant c . In particular, triangles are objects of constant complexity. The boundary of an object consists of a number of *edges*, which are connected portions of algebraic curves, and *vertices*, where these curves join together. To simplify our presentation, in the remainder of the paper we will disregard the possibility that an object can have no vertices: this can happen, if the object is, for example, a circular disk; we then introduce an artificial vertex, say at the highest point of the object.

The union $\mathcal{U}(\mathcal{F})$ of a set \mathcal{F} of objects in the plane is a planar set bounded by a set of *vertices*, which include original object vertices and union boundary points where edges of different objects intersect, and *edges*, each being a maximal connected portion of an object edge appearing on the union boundary and not containing a vertex in its relative interior. For the remainder of this paper, the *complexity* $|\mathcal{U}(\mathcal{F})|$ of $\mathcal{U}(\mathcal{F})$ is defined as the number of union *vertices*.

We aim to bound the maximum complexity of the union \mathcal{F} of n fat objects in the plane, where the objects are either constant description complexity locally γ -fat objects or γ -fat triangles. We assume from now on that the objects of \mathcal{F} are in *general position*, meaning that no object vertex appears on the boundary of another object, no three object boundaries meet at a common point, and every intersection of two object boundaries is a discrete set of points, each of which is a proper crossing of the two boundaries. A standard perturbation argument [34, Chapter 7] shows that the maximum union complexity, for fixed n, c, γ , is achieved, up to a constant factor, by sets in general position.

Let s be the maximum number of times any two edges of two different objects intersect. Then s must be a constant due to our assumption that the sets in \mathcal{F} have constant complexity, and that any pair of boundaries intersect in a discrete set.⁴

⁴Concretely, this is a consequence of Bézout's theorem [25].

Our proof, as most previous proofs on the union complexity of fat objects, proceeds roughly as follows. We cover the boundaries of the objects in \mathcal{F} by simpler objects, contained within the original objects, which are then grouped into “canonical” families. The net effect is that vertices of the union of \mathcal{F} appear as vertices of the union of the new objects, so bounding the complexity of this latter kind of unions is sufficient. Within each family the objects have a certain “canonical” shape; this allows us to prove a good bound on the union complexity of each family separately. It then remains to derive a bound on the overall union from the bounds for the individual families. Combining the unions of the subfamilies can be done using existing tools, which we briefly review.

We say that a function f defined on non-negative integers is *superadditive* if $f(n) + f(m) \leq f(n + m)$, for all $n, m \geq 0$.

Observation 2.1. *Consider any class \mathcal{C} of bounded planar objects of constant complexity, closed under translation. Let $U_{\mathcal{C}}(n)$ be the maximum complexity of the union of any n objects from \mathcal{C} . Then $U_{\mathcal{C}}$ is superadditive.*

Proof. Let \mathcal{U}_1 be a union of n objects from \mathcal{C} achieving $|\mathcal{U}_1| = U_{\mathcal{C}}(n)$. Define \mathcal{U}_2 similarly, for m objects. We can assume that $\mathcal{U}_1 \cap \mathcal{U}_2 = \emptyset$; indeed, if this is not the case then the translation invariance of \mathcal{C} and the boundedness of the objects ensure that we can move \mathcal{U}_2 so that it becomes disjoint from \mathcal{U}_1 . Hence, $U_{\mathcal{C}}(n + m) \geq |\mathcal{U}_1 \cup \mathcal{U}_2| = |\mathcal{U}_1| + |\mathcal{U}_2| = U_{\mathcal{C}}(n) + U_{\mathcal{C}}(m)$, as claimed. \square

Combining unions of triangles. In what follows, we denote the number of *holes* in the union $\mathcal{U}(\mathcal{F})$ (that is, connected components of the complement of the union) by $H(\mathcal{F})$. Edelsbrunner *et al.* [15] have shown that the complexity of the union $\mathcal{U}(\mathcal{F}_1 \cup \mathcal{F}_2)$ of two sets \mathcal{F}_1 and \mathcal{F}_2 of triangles is bounded by the sum of the individual union complexities $|\mathcal{U}(\mathcal{F}_1)|$ and $|\mathcal{U}(\mathcal{F}_2)|$, and an overhead term which is linear in the total number of triangles and the number of holes in the combined union. This is made precise in the so-called Combination Lemma stated below.

We will sometimes need to apply the Combination Lemma within some constant-complexity “window” R , such as a triangle or a trapezoid, and then the complexities of the three unions and the number of holes will refer to the portion of the configuration clipped within R . It is easy to check that the lemma continues to hold in this context, as the number of additional holes caused by the clipping is only $O(n)$. We denote the union complexity of a set \mathcal{F} of objects inside a region R by $|\mathcal{U}(\mathcal{F}) \cap R|$. In other words, $|\mathcal{U}(\mathcal{F}) \cap R| = |\mathcal{U}(\{o \cap R : o \in \mathcal{F}\})|$. The number of holes defined by the objects in \mathcal{F} inside R —more precisely, the number of holes of $\mathcal{F} \cup \{\partial R\}$ inside R —is denoted by $H_R(\mathcal{F})$. The (modified version of the) Combination Lemma can now be stated as follows.

Combination Lemma (see Edelsbrunner *et al.* [15]). *Let \mathcal{F}_1 and \mathcal{F}_2 be two families of a total of n triangles in the plane, and let R be a polygonal region of constant complexity. Then we have*

$$|\mathcal{U}(\mathcal{F}_1 \cup \mathcal{F}_2) \cap R| \leq |\mathcal{U}(\mathcal{F}_1) \cap R| + |\mathcal{U}(\mathcal{F}_2) \cap R| + O(n + H_R(\mathcal{F}_1 \cup \mathcal{F}_2)).$$

Proof. Let R_{out} be a large rectangle that contains $\mathcal{U}(\mathcal{F}_1 \cup \mathcal{F}_2)$ and R in its interior. Let T be a collection of (a constant number of) triangles forming a triangulation of $R_{\text{out}} \setminus R$. Then $|\mathcal{U}(\mathcal{F}_1 \cup \mathcal{F}_2) \cap R| = |\mathcal{U}(\mathcal{F}_1 \cup \mathcal{F}_2 \cup T)| - 4$ (the four vertices of R are not counted in the left-hand side) and $H_R(\mathcal{F}_1 \cup \mathcal{F}_2) = H(\mathcal{F}_1 \cup \mathcal{F}_2 \cup T) - 1$ (since by definition H counts the “unbounded hole” and H_R does not). Hence, using the regular Combination Lemma—the version not involving a

window R —we can derive

$$\begin{aligned}
|\mathcal{U}(\mathcal{F}_1 \cup \mathcal{F}_2) \cap R| &\leq |\mathcal{U}(\mathcal{F}_1 \cup \mathcal{F}_2 \cup T)| \\
&\leq |\mathcal{U}(\mathcal{F}_1 \cup T)| + |\mathcal{U}(\mathcal{F}_2 \cup T)| + O(n + H(\mathcal{F}_1 \cup \mathcal{F}_2 \cup T)) \\
&\leq |\mathcal{U}(\mathcal{F}_1) \cap R| + |\mathcal{U}(\mathcal{F}_2) \cap R| + O(n + H_R(\mathcal{F}_1 \cup \mathcal{F}_2)), \quad \square
\end{aligned}$$

where in the last inequality the extra additive constants are subsumed in the term $O(n)$.

Notice that the Combination Lemma does not assume fatness, and thus can be applied to arbitrary triangles. (In fact, the formulation in [15] is even more general, and is stated for arbitrary arrangements of segments and any subset of cells in the overlay of two sub-arrangements.) Notice also the crucial property that the terms $|\mathcal{U}(\mathcal{F}_1) \cap R|$ and $|\mathcal{U}(\mathcal{F}_2) \cap R|$ appear with coefficient 1. This is not the case in the more general version of the lemma, which applies to curved objects [34] (nor is it the case in the Merging Lemma below) and this makes the analysis of the unions of curved objects more involved, and somewhat less tight.

When applying the Combination Lemma, the main issue is to bound the number of holes in the combined union. For fat triangles, it is known that this number is linear [31, 32]. The lemma below makes this precise, including the improved dependence on γ as obtained by Pach and Tardos [32]. Again, we will need the lemma inside a window R . Obviously it also holds in that setting, since, as already noted, the window cannot increase the number of holes by more than $O(n)$.

Few Holes Lemma (Pach and Tardos [32]). *Let \mathcal{F} be a collection of n γ -fat triangles in the plane and let R be a polygonal region of constant complexity. Then $H_R(\mathcal{F}) = O(\frac{n}{\gamma} \log \frac{1}{\gamma})$. This bound is tight in the worst case, up to a factor of $\log \frac{1}{\gamma}$.*

Combining unions of fat curved objects. The Combination Lemma, in its strong form stated above, does not apply to curved objects. Hence, we need a different tool when we want to combine unions of fat curved objects. This is supplied by the so-called Merging Lemma [6]. As in the case of triangles, we will sometimes need to bound the union complexity of a set of fat objects restricted to some (in this case possibly curved) constant-complexity region R .

Merging Lemma (De Berg [6]). *Let \mathcal{F}_1 and \mathcal{F}_2 be two families of locally γ -fat objects, and let R be a region in the plane of constant complexity. Then we have*

$$|\mathcal{U}(\mathcal{F}_1 \cup \mathcal{F}_2) \cap R| = O(1/\gamma) \cdot (|\mathcal{U}(\mathcal{F}_1) \cap R| + |\mathcal{U}(\mathcal{F}_2) \cap R|),$$

where the constant of proportionality also depends on the complexity of the individual objects of $\mathcal{F}_1 \cup \mathcal{F}_2$.

Note that, unlike the Combination Lemma, the Merging Lemma requires the objects in \mathcal{F}_1 and \mathcal{F}_2 to be fat. Another important difference is that the terms $|\mathcal{U}(\mathcal{F}_1)|$ and $|\mathcal{U}(\mathcal{F}_2)|$ appear in the above bound with a coefficient $O(1/\gamma)$. When we use the Merging Lemma only once this is not a major issue, but in repeated applications of the lemma in a recursive scheme the constant factor can blow up. This is the main reason why our bound for curved objects is weaker than the bound we obtain for triangles.

Proof overview. Our proof proceeds as follows.

First, we cover each object in \mathcal{F} by so-called *towers*; see Figure 4(b) for an illustration. (In fact, the towers need not cover the entire object: it is sufficient that they cover the boundary of each object, so that the union boundary of the original objects is a subset of the union boundary of all the towers.) The collection of towers used in the cover can be partitioned into classes, such that the towers in each class have a certain canonical shape. The details of the reduction to towers are described in Section 3.

In Section 4 we then proceed to bound the union complexity of each of the classes, using a recursive counting scheme. One of the ingredients needed to apply this technique is a sharp bound on the union complexity of fat objects stabbed by a common vertical line. Such a bound was already known for fat (canonical) triangles [31]—see also the recent improvements of Pettie [33]—but not for curved fat objects. Thus, a main step in the proof will be to obtain such a bound for towers. In particular, we will prove that the union complexity of n towers stabbed by a vertical line is $O(\lambda_{s+2}(n))$, where s is the maximum number of intersections between a pair of tower diagonals. This bound improves to $O(n)$ for the case of polygonal towers (where $s = 1$). Using our recursive counting scheme we then obtain a bound on the union complexity of towers. Note that the improved bound for polygonal towers stabbed by a vertical line is better than the known bound on the union complexity of fat canonical triangles stabbed by a line [31, 33]. The improvement also carries over to the final bound we obtain after applying the recursive counting scheme.

After bounding the union complexity of a set of towers from a fixed class, it remains to combine the different classes of towers using the Combination Lemma (in the case of triangles) or the Merging Lemma (in the case of curved objects). This gives us our final bounds on the union complexity of fat objects: $n \cdot 2^{O(\log^* n)}$ for locally γ -fat objects and $O(n \log^* n)$ for γ -fat triangles; see Theorems 4.10 and 4.11, respectively (where the dependence of these bounds on γ is also specified).

3 Reduction to towers

We start by reducing the problem of bounding the union complexity of \mathcal{F} to that of bounding the union complexity of a set of objects with a more restricted shape, which we call *towers*. This is done in two steps: first we cover the boundary of each object in \mathcal{F} by a collection of “well-adjusted” objects with a triangular shape, and then we cover those intermediate objects by towers. Along the way we eliminate the dependence of the analysis on the fatness parameter γ by reducing the problem to subproblems, each involving objects of fatness bounded by an absolute constant; the fatness parameter will then only affect the number of subproblems we have to deal with and (for curved objects) the overhead involved in combining subproblems via the Merging Lemma.

Step 1: Covering objects in \mathcal{F} by well-adjusted (quasi-)triangles

We first describe how to handle γ -fat triangles and then move on to locally γ -fat objects.

Covering boundaries of fat triangles by well-adjusted triangles. Our approach is similar to the cover used by Matoušek *et al.* [31], though we modify the construction to make it more consistent with that for the curved case.

Set $\gamma^* := \gamma/3$, and let \mathcal{D} be a collection of $\lceil 2\pi/\gamma^* \rceil = \Theta(1/\gamma)$ evenly spaced canonical orientations such that the angle between any two consecutive orientations⁵ is at most γ^* . (Since $\gamma \leq \pi/3$,

⁵We do not identify opposite orientations.

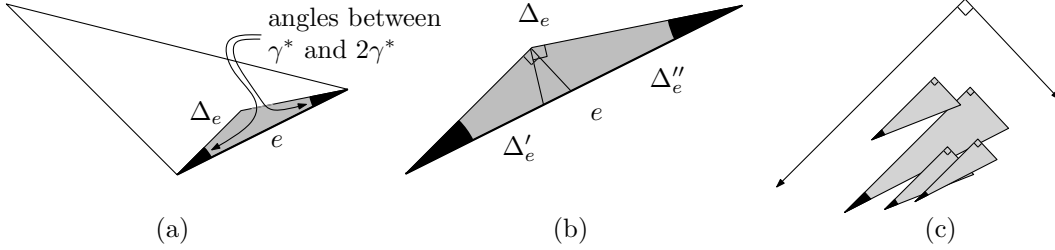


Figure 2: (a) Covering an edge e of a fat triangle by a semi-canonical triangle Δ_e . (b) Replacing Δ_e by two semi-canonical right triangles. (c) A single family of semi-canonical right triangles (which includes Δ'_e). All triangles in this family have their legs parallel to two indicated orientations.

we have $\gamma^* \leq \pi/9$ and $|\mathcal{D}| \geq 18$.) Let $\Delta \in \mathcal{F}$ be a γ -fat triangle. We cover each edge e of Δ with a triangle $\Delta_e \subset \Delta$ such that (i) e is an edge of Δ_e , (ii) the other two edges have orientations in \mathcal{D} and are contained in Δ , and (iii) the angles of Δ_e at the endpoints of e are at least γ^* and at most $2\gamma^*$. We call such a triangle *semi-canonical*; see Figure 2(a). We next cover Δ_e with two right triangles (refer to Figure 2(b)) Δ'_e and Δ''_e with $\Delta_e = \Delta'_e \cup \Delta''_e$. Each of the two triangles has its right-angle vertex at the vertex of Δ_e opposite to e and one of its other vertices at a respective endpoint of e . Since $\gamma \leq \pi/3$, the smallest angle of the subtriangle is at the endpoint of e and, as we recall, it is between γ^* and $2\gamma^*$. Applying this procedure to every edge of each triangle $\Delta \in \mathcal{F}$, we obtain a family \mathcal{G} of $6n$ “semi-canonical” triangles, with $|\mathcal{U}(\mathcal{F})| \leq |\mathcal{U}(\mathcal{G})|$. The family \mathcal{G} can be partitioned into $O(1/\gamma)$ subfamilies \mathcal{G}_i such that the triangles from each subfamily are right-angle, have a fixed orientation from \mathcal{D} for their longer leg⁶ and an orthogonal orientation for the other leg—see Figure 2(c).

Now consider one such subfamily \mathcal{G}_i . It consists of right triangles so that, without loss of generality, their right-angle vertex is the lower-left vertex, their legs are vertical and horizontal, and their smallest angle is at their top vertex, measuring between γ^* and $2\gamma^*$; hence their diagonals have slopes of absolute value between $\cot 2\gamma^*$ and $\cot \gamma^*$. We now scale all triangles of \mathcal{G}_i vertically by a factor of $\tan \gamma^*$, that is, we apply the transformation $\Gamma: (x, y) \mapsto (x, y \tan \gamma^*)$. The resulting right triangles all have a diagonal with slope between -1 and $-\frac{\tan \gamma^*}{\tan 2\gamma^*}$. Since $\gamma^* = \gamma/3 \leq \pi/9$, we have

$$\frac{\tan \gamma^*}{\tan 2\gamma^*} = \frac{\tan \gamma^*}{2 \tan \gamma^* / (1 - \tan^2 \gamma^*)} = \frac{1 - \tan^2 \gamma^*}{2} \geq \frac{1 - \tan^2(\pi/9)}{2} > \frac{1}{3},$$

so the slopes lie in $[-1, -\frac{1}{3})$. We call such (right) triangles *well-adjusted*.

Clearly, the union complexity of a single family \mathcal{G}_i is not affected by the scaling transformation. The following lemma summarizes how bounding the complexity of a collection of well-adjusted triangles yields a similar bound for the union complexity of any set of γ -fat triangles, for any fixed $0 < \gamma \leq \pi/3$.

Lemma 3.1. *Let \mathcal{F} be a set of n γ -fat triangles. Then there exists a collection of $O(1/\gamma)$ families \mathcal{G}_i , each consisting of $(\gamma/3)$ -fat right triangles, such that the following holds:*

(i) $|\mathcal{U}(\mathcal{F})| \leq |\bigcup_i \mathcal{U}(\mathcal{G}_i)|$.

⁶The *legs* of a right triangle are the two sides incident to the right-angle vertex; we refer to the third side as the *diagonal* rather than *hypotenuse* to be consistent with our terminology in the curved case studied below.

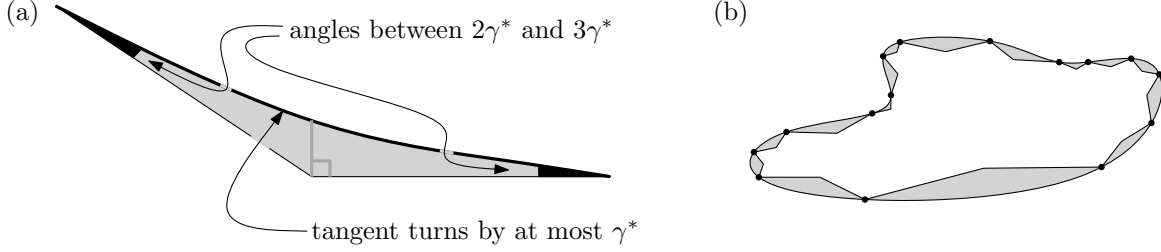


Figure 3: (a) A γ^* -quasi-triangle. (b) Covering the boundary of a curved object by semi-canonical quasi-triangles.

(ii) Each family \mathcal{G}_i can be transformed (without changing the combinatorial structure of the union) into a collection of well-adjusted triangles, meaning that the transformed triangles have

- a vertical leg and a horizontal leg, meeting at its lower-left vertex,
- a diagonal with slope in the range $[-1, -\frac{1}{3}]$.

(iii) $\sum_i |\mathcal{G}_i| = 6n$.

Covering locally fat objects with well-adjusted quasi-triangles. We now describe a procedure, similar to the one above, for covering locally fat (and possibly curved) objects instead of just triangles.

We start by covering the objects by certain semi-canonical shapes (the number of such shapes, for each object, is a constant that depends on the complexity of this object). The shapes we shall use are so-called quasi-triangles, as defined next and illustrated in Figure 3(a). A γ^* -quasi-triangle Δ is a triangle-like object with the following properties. It is delimited by two straight line segments and an arc. The arc is smooth, has no inflection points, and its tangent turns by an angle of at most γ^* along the arc. The angle between the two straight edges of Δ is between $\pi - 7\gamma^*$ and $\pi - 3\gamma^*$, while the angles at the other two corners of Δ (defined as the angles between the straight edges and the tangents to the curved edge) are between $2\gamma^*$ and $3\gamma^*$. The γ^* -quasi-triangle illustrated in Figure 3 is *concave*, in the sense that, as we traverse its boundary in counterclockwise direction, the curved edge turns clockwise. Otherwise, the γ^* -quasi-triangle is *convex*; in this latter case it is indeed a convex set.

Now let \mathcal{F} be a set of n locally γ -fat objects of bounded description complexity. Let⁷ $\gamma^* = \gamma\pi/20$. As before, we let \mathcal{D} be a collection of $\lceil 2\pi/\gamma^* \rceil = \Theta(1/\gamma)$ evenly spaced canonical orientations such that the angle between any two consecutive directions in \mathcal{D} is at most γ^* . We call a γ^* -quasi-triangle *semi-canonical* if its two straight edges have orientations in \mathcal{D} . The following property has been shown by De Berg [6, Section 4] (see also Figure 3(b) for an illustration):

Lemma 3.2 (De Berg [6]). *One can cover the boundaries of the objects in \mathcal{F} by a family \mathcal{G} of γ^* -quasi-triangles as above,⁸ so that (i) each quasi-triangle is contained in an original object,*

⁷Recall that the parameter γ appearing in locally fat objects has a somewhat different meaning than that for fat triangles.

⁸The definition in [6], where the term γ -standard quasi-triangle was used, is slightly different: the angle between the two straight edges can be between $\pi - 7\gamma^*$ and $\pi - \gamma^*$, and the angles at the two other corners are only required to be at least γ^* . However, the latter angles are actually proved to be between $2\gamma^*$ and $3\gamma^*$ [6, Lemma 4.3]. Together with the fact that the tangent does not turn by more than γ^* this implies that the angle between the two straight edges cannot be greater than $\pi - 3\gamma^*$.

(ii) a total of $O(n/\gamma^2)$ quasi-triangles are required (with a constant of proportionality that depends on the object complexity), and (iii) the quasi-triangles are semi-canonical.

Property (iii) implies that the quasi-triangles can be grouped into $O(1/\gamma)$ subfamilies \mathcal{G}_i , each of which consists of quasi-triangles whose straight edges have two fixed orientations (which are close to each other in \mathcal{D}).

Consider one such subfamily \mathcal{G}_i . We further cover each γ^* -quasi-triangle $\Delta \in \mathcal{G}_i$ with two right-angle shapes, called *right γ^* -quasi-triangles*. Each such shape is bounded by two straight edges (its *legs*) and a curved edge (its *diagonal*), so that (i) the legs are at right angles to each other, (ii) the smallest of the two other angles (defined as above) is between $2\gamma^*$ and $3\gamma^*$, (iii) the third angle is between $\pi/2 - \gamma^*$ and $\pi/2 - 4\gamma^*$, and (iv) the diagonal has no inflection points and turns by at most γ^* . In Figure 3(a) one of these shapes is indicated, namely the one delimited by the bottom edge of Δ , part of its diagonal, and a segment (shown in dark gray) extending from the lower left corner orthogonally to the bottom edge. The other shape needed to cover Δ is delimited by the left edge, a part of the diagonal, and a segment (not shown in the figure) extending from the lower left corner orthogonally to the left edge. The following lemma was proven⁹ in an earlier paper of De Berg [7, proof of Lemma A.1].

Lemma 3.3 (De Berg [7]). *A right γ^* -quasi-triangle is locally $\Theta(\gamma^*)$ -fat.*

To summarize: we can cover the boundaries of the objects of \mathcal{F} by a total of $O(n/\gamma^2)$ right γ^* -quasi-triangles, for a suitable $\gamma^* = \Theta(\gamma)$; these quasi-triangles are locally fat and can be grouped into $O(1/\gamma)$ semi-canonical families \mathcal{G}_i . As in the case of triangles, $|\mathcal{U}(\mathcal{F})|$ is upper bounded by the union complexity of the new collection $\mathcal{G} := \bigcup_i \mathcal{G}_i$.

Similar to the case of triangles, we now want to make the quasi-triangles well-adjusted. To this end, consider a single semi-canonical family \mathcal{G}_i of right γ^* -quasi-triangles. Without loss of generality, we may assume that each quasi-triangle in \mathcal{G}_i is such that (i) its legs are vertical and horizontal, (ii) its right-angle vertex is the lower-left vertex, and (iii) its smallest angle is at the top vertex. Analogously to the case of “straight” triangles, we apply a vertical scaling $\Gamma: (x, y) \mapsto (x, y \tan 2\gamma^*)$ to the entire family, producing a set of what we call *well-adjusted quasi-triangles*. Let us examine the absolute value of the slope of (the tangent to) the diagonal at any point. Since the diagonal has no inflection points, it is at its steepest at one of its endpoints and at its shallowest at the other. Before scaling, the angle with vertical ranged between γ^* and $4\gamma^*$, so the absolute value of the slope ranged between $\cot 4\gamma^*$ and $\cot \gamma^*$. Hence the slope of the diagonal in a well-adjusted triangle lies in the range $[-\frac{\tan 2\gamma^*}{\tan \gamma^*}, -\frac{\tan 2\gamma^*}{\tan 4\gamma^*}]$. Recall that $0 < \gamma \leq 1/4$ and that $\gamma^* = \gamma\pi/20$. Hence, γ^* ranges between 0 and $\pi/80$, which implies that

$$\frac{\tan 2\gamma^*}{\tan 4\gamma^*} = \frac{1 - \tan^2 2\gamma^*}{2} \geq \frac{1 - \tan^2(\pi/40)}{2} \geq 0.496,$$

and

$$\frac{\tan 2\gamma^*}{\tan \gamma^*} = \frac{2}{1 - \tan^2 \gamma^*} \leq \frac{2}{1 - \tan^2(\pi/80)} \leq 2.0031.$$

⁹In that paper the right quasi-triangle was obtained in a slightly different manner, namely by applying a shear transformation to a γ^* -quasi-triangle rather than covering it by two right-angle pieces, but the fatness proof still goes through. (Using the shear transformation to make the quasi-triangles right-angled does not work for us, because then the next step—making them well-adjusted—is problematic.)

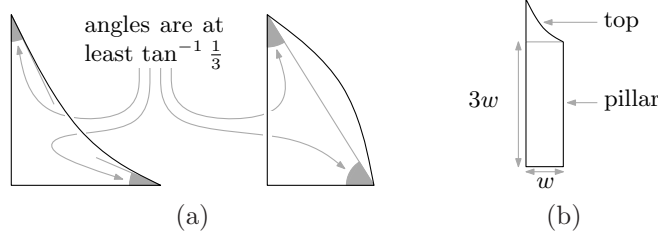


Figure 4: (a) Two well-adjusted quasi-triangles. (b) A tower.

Thus we can safely assume that the slope of the diagonal always stays within $(-3, -\frac{1}{3})$.

Recall that we assumed that any pair of boundary arcs of the original objects in \mathcal{F} intersect in at most s points. Obviously this carries over to the diagonals of our well-adjusted quasi-triangles. In addition, recall the properties of a γ^* -quasi-triangle: its diagonal arc has no inflection points, its tangent turns by an angle of at most γ^* , the angle between the two straight edges is between $\pi - 7\gamma^*$ and $\pi - 3\gamma^*$, and the angles at the other two corners are between $2\gamma^*$ and $3\gamma^*$. This implies that the diagonal arc is a convex or concave x - and y -monotone, as it cannot become parallel to one of the legs. This property too carries over to the diagonals of our well-adjusted right quasi-triangles (as shown above, their slopes lie in $(-3, -\frac{1}{3})$). The following lemma summarizes the discussion above (and, in particular, the properties stated in Lemma 3.2).

Lemma 3.4. *Let \mathcal{F} be a set of n locally γ -fat objects. Then there exists a collection of $O(1/\gamma)$ families \mathcal{G}_i , each consisting of right γ^* -quasi-triangles for a suitable $\gamma^* = \Theta(\gamma)$, such that the following holds:*

- (i) $|\mathcal{U}(\mathcal{F})| \leq |\bigcup_i \mathcal{U}(\mathcal{G}_i)|$.
- (ii) *Each family \mathcal{G}_i can be transformed (without changing the combinatorial structure of its union) into a collection of well-adjusted quasi-triangles, meaning that the transformed quasi-triangles have*
 - *a vertical leg and a horizontal leg, meeting at its lower-left vertex,*
 - *a convex or concave x - and y -monotone arc (without inflection points) as diagonal, and the tangent along the diagonal has a slope whose value stays in the range $(-3, -\frac{1}{3})$.*

Moreover, the diagonals of any pair of quasi-triangles intersect in at most s points.

- (iii) $\sum_i |\mathcal{G}_i| = O(n/\gamma^2)$.

Step 2: Covering semi-canonical (quasi-)triangles by towers

After applying the reduction described above, we are left with the problem of bounding the union complexity of a semi-canonical set of well-adjusted right (quasi-)triangles. To do so, we cover these (quasi-)triangles by so-called *towers* (and by some fat rectangles), as explained next. Most of the remainder of the paper will then be devoted to bounding the union complexity of a set of towers, whereas at the end of this section we show how to derive the bounds on the union complexity of the original objects (that is, γ -fat triangles or locally γ -fat objects) given the union complexity of towers—see Lemmas 3.8 and 3.9 for more details.

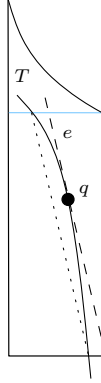


Figure 5: Proof of Observation 3.5.

We start by defining towers. Define the *aspect ratio* of an axis-parallel rectangle to be its height (y -span) divided by its width (x -span). A *vertical tower* is the union of a rectangle of aspect ratio 3 and a well-adjusted quasi-triangle (of the canonical form asserted in Lemma 3.4) whose bottom edge coincides with the top edge of the rectangle—see Figure 4(b) for an illustration. We call the rectangular part of a tower T its *pillar* and the quasi-triangular part its *top*. We denote these parts by $\text{pillar}(T)$ and $\text{top}(T)$, respectively. A *horizontal tower* is defined similarly, except that now the pillar has aspect ratio $1/3$ and the right edge of the pillar coincides with the left edge of the quasi-triangular “top” (clearly, a misnomer, but we keep this terminology for the sake of uniformity), whose right angle is at the bottom endpoint of that edge. These definitions also apply in the triangular case: there the top of a tower is simply a well-adjusted right triangle with one horizontal edge and one vertical edge, with its right angle at its lower left vertex. Where appropriate we will refer to towers that arise in the case of triangles as *straight towers*, and to towers that arise in the case of locally fat objects as *curved towers*.

Observation 3.5. *Let T be a vertical tower and e be the diagonal edge of another vertical tower. Then e cannot cross both the bottom and the top edges of the pillar of T . An analogous statement holds for the left and right edges of pillars of horizontal towers.*

Proof. If the smooth descending curve e crossed both the top and bottom edges of $\text{pillar}(T)$, which is a rectangle of aspect ratio 3, e would have to contain a point q so that the tangent to e at q has slope at most -3 (see Figure 5), as follows from the Mean Value Theorem. However, by construction, the slope of any such tangent is strictly larger than -3 (this holds both for triangles and for quasi-triangles), which results in a contradiction. Horizontal pillars are handled in a symmetric manner, using the upper bound of $-1/3$ on the slopes of their tangents. \square

We will also need the following property.

Lemma 3.6. *A tower is locally $\Theta(1)$ -fat.*

Proof. Let T be a tower. By applying Lemma 3.3 with $\gamma^* = \Theta(1)$, we see that $\text{top}(T)$ is locally $\Theta(1)$ -fat. Obviously, $\text{pillar}(T)$ is locally $\Theta(1)$ -fat as well. Now consider a disk D with center $p \in T$ that does not fully contain T in its interior. We must show that $\text{area}(D \cap T)/\text{area}(D) = \Omega(1)$.

Assume that $p \in \text{top}(T)$; a similar argument can be used when $p \in \text{pillar}(T)$. If $\text{top}(T)$ is not contained in the interior of D , then $\text{area}(D \cap T)/\text{area}(D) = \Omega(1)$ because $D \cap \text{top}(T) \subseteq D \cap T$ and $\text{top}(T)$ is locally $\Theta(1)$ -fat. If it is fully contained in the interior of D , then consider the largest disk D' with center p whose boundary meets $\text{top}(T)$. It is easy to see that the radius of D cannot be more than a constant factor larger than that of D' (or else D would have “swallowed” the entire tower). Hence,

$$\begin{aligned} \text{area}(D \cap T)/\text{area}(D) &\geq \text{area}(D' \cap T)/\text{area}(D) \\ &= \Omega(\text{area}(D' \cap T)/\text{area}(D')) && \text{(because } \text{area}(D)/\text{area}(D') = \Theta(1)\text{)} \\ &= \Omega(\text{area}(D' \cap \text{top}(T))/\text{area}(D')) \\ &= \Omega(1) && \text{(because } \text{top}(T) \text{ is locally } \Theta(1)\text{-fat),} \end{aligned}$$

thus finishing the proof. \square

Next we show that any well-adjusted (quasi-)triangle Δ can be covered by three vertical towers, three horizontal towers, and six fat rectangles, all contained in Δ . Assume, without loss of generality, that the longer of the vertical and the horizontal edges of Δ , say, the vertical edge, has length 1. Let σ be the largest square contained in Δ whose lower-left corner coincides with the lower-left corner of Δ ; this square is shown in Figure 6(a).

Claim 3.7. *The side length of σ is at least $1/6$.*

Proof. Let x denote the side length of σ . Assume first that Δ is a concave quasi-triangle. In this case Δ is contained in the union of σ and two right triangles; one of these triangles has a vertical edge of length $1 - x$ and a horizontal edge of length x , the other has a vertical edge of length x and a horizontal edge of length at most $1 - x$; see Figure 6(a). Hence,

$$\text{area}(\Delta) \leq x^2 + 2 \cdot \frac{x(1-x)}{2} = x.$$

On the other hand, since the absolute value of the slope of the diagonal at the top vertex of Δ is at most 3, we have

$$\text{area}(\Delta) \geq \frac{1}{6},$$

thus proving the claim. Note that the argument applies verbatim if Δ is a triangle. If Δ is a convex quasi-triangle, shrink it to a triangle Δ' whose diagonal is the segment connecting the top and right vertices of Δ . Note that the absolute value of its slope is at most 3. The largest square with the above properties which is inscribed in Δ is larger than the largest such square inscribed in Δ' . The preceding analysis shows that the side length of the latter square is larger than $1/6$, so the same holds for the former square, thus establishing the claim for the convex case too. \square

We now partition σ by two vertical cuts into three rectangles whose aspect ratio is 3 and extend these rectangles upward until they reach the diagonal edge, as shown in Figure 6(b). This partitions the left part of Δ —that is, the part of Δ to the left of the right edge of σ —into three vertical slices. These slices are not yet proper towers, because (except for the rightmost slice) their pillars have aspect ratio greater than 3. Hence, we further partition each of the two leftmost slices into a tower and a rectangular remainder, as illustrated in Figure 6(c). Note that the remainders are

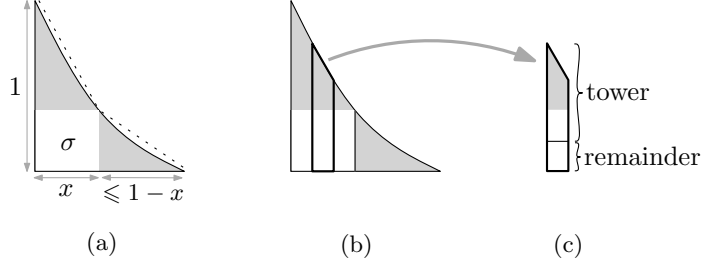


Figure 6: (a) The square σ and the two (shaded) right quasi-triangles that together form Δ . (b) Partitioning the left part of Δ into three slices. (c) Partitioning a slice into a tower and a rectangular remainder.

axis-aligned rectangles of constant aspect ratio: their horizontal edges have length at least $1/18$ (and at most $1/3$) and their vertical edges have length less than 1 (and at least some constant that depends on the slope of the diagonal edge of Δ). Note also that the remainders can “stick out” of σ .

We have thus covered the portion of Δ that lies to the left of σ 's right edge by three vertical towers and three fat rectangles. We can cover the part of Δ below the top edge of σ in a similar manner, using horizontal towers and fat rectangles. Together we obtain a cover of Δ by six towers (three vertical and three horizontal ones) and six axis-aligned rectangles of constant aspect ratio.

This finishes the description of our reduction to towers. The following two lemmas summarize the relation between our original problem—bounding the union complexity of sets of fat triangles or sets of locally-fat objects—and the union complexity of towers.

Lemma 3.8. *Let $U(m)$ denote the superadditive function measuring the maximum union complexity of any set of m vertical or m horizontal straight towers. Then the union complexity of any set \mathcal{F} of n γ -fat triangles is at most $U(36n) + O(\frac{n}{\gamma} \log^2 \frac{1}{\gamma})$.*

Proof. We first relate the union complexity of towers to the union complexity of well-adjusted triangles. To this end, let T be a set of n well-adjusted triangles. We cover each well-adjusted triangle by three horizontal towers, three vertical towers, and six axis-aligned rectangles of constant aspect ratio, as explained above. Let T_{hor} and T_{vert} be the resulting sets of horizontal and vertical towers, respectively, and let A be the resulting set of rectangles. Note that any tower can be covered by a constant number of triangles of constant fatness. The Few Holes Lemma thus implies that $\mathcal{U}(T_{\text{hor}} \cup T_{\text{vert}})$ has $O(n)$ holes (with an absolute constant of proportionality). Using the Combination Lemma, we obtain

$$\begin{aligned}
 |\mathcal{U}(T_{\text{hor}} \cup T_{\text{vert}})| &\leq |\mathcal{U}(T_{\text{hor}})| + |\mathcal{U}(T_{\text{vert}})| + O(n) \\
 &\leq U(3n) + U(3n) + O(n) \\
 &\leq U(6n) + O(n),
 \end{aligned}$$

where the last step uses the superadditivity of $U(\cdot)$. Now, since each rectangle in A is axis-aligned and has constant aspect ratio, we have $|\mathcal{U}(A)| = O(n)$ (this is because each such rectangle can be covered by a constant number of squares, and their union complexity is known to be linear [28]). In addition, each rectangle in A can be covered by two triangles of constant fatness, so $\mathcal{U}((T_{\text{hor}} \cup$

$T_{\text{vert}} \cup A$) has $O(n)$ holes (again, with an absolute constant of proportionality) by the Few Holes Lemma. Hence, by the Combination Lemma we get

$$|\mathcal{U}((T_{\text{hor}} \cup T_{\text{vert}}) \cup A)| \leq |\mathcal{U}(T_{\text{hor}} \cup T_{\text{vert}})| + |\mathcal{U}(A)| + O(n) \leq U(6n) + O(n).$$

Therefore, denoting by $U_{\text{w}}(n)$ the maximum union complexity of n well-adjusted triangles, we have $U_{\text{w}}(n) \leq U(6n) + O(n)$.

Now that we have a bound for well-adjusted triangles, we can proceed to arbitrary fat triangles. Let \mathcal{F} be a set of n γ -fat triangles. By Lemma 3.1 we know that $|\mathcal{U}(\mathcal{F})| \leq |\mathcal{U}(\mathcal{G}_1 \cup \dots \cup \mathcal{G}_\kappa)|$, where each \mathcal{G}_i is a collection of $(\gamma/3)$ -fat right triangles, with $\kappa = O(1/\gamma)$ and $\sum_i |\mathcal{G}_i| = 6n$. Each collection \mathcal{G}_i can be transformed to a collection of well-adjusted triangles, so the union complexity of \mathcal{G}_i is at most $U_{\text{w}}(n_i) \leq U(6n_i) + O(n_i)$, where $n_i := |\mathcal{G}_i|$.

We combine the unions of the subfamilies \mathcal{G}_i in a tree-like manner, as follows. We construct a balanced binary tree of depth $O(\log \kappa) = O(\log \frac{1}{\gamma})$ whose leaves represent the subfamilies \mathcal{G}_i . Each internal node ν represents the union $\mathcal{G}(\nu)$ of the subfamilies of its children, so that the root represents $\mathcal{G}_1 \cup \dots \cup \mathcal{G}_\kappa$. For an internal node ν , put $n_\nu := |\mathcal{G}(\nu)|$, and let ν_1 and ν_2 denote the two children of ν . Using the Combination Lemma, the Few Holes Lemma, and the fact that the triangles in each \mathcal{G}_i are $(\gamma/3)$ -fat, we see that

$$|\mathcal{U}(\mathcal{G}(\nu))| \leq |\mathcal{U}(\mathcal{G}(\nu_1))| + |\mathcal{U}(\mathcal{G}(\nu_2))| + O\left(\frac{n_\nu}{\gamma} \log \frac{1}{\gamma}\right).$$

In other words, the union complexity at a node ν is at most the sum of the union complexities at its children, plus an additive overhead term $O(\frac{n_\nu}{\gamma} \log \frac{1}{\gamma})$. This means that the union complexity at the root of the tree is bounded by

$$\sum_{i=1}^{\kappa} U_{\text{w}}(n_i) + \sum_{\nu} O\left(\frac{n_\nu}{\gamma} \log \frac{1}{\gamma}\right) \leq \sum_{i=1}^{\kappa} (U(6n_i) + O(n_i)) + \sum_{\nu} O\left(\frac{n_\nu}{\gamma} \log \frac{1}{\gamma}\right),$$

which is at most $U(36n) + O(\frac{n}{\gamma} \log^2 \frac{1}{\gamma})$, because $U(\cdot)$ is superadditive, $\sum_{i=1}^{\kappa} n_i = 6n$, and $\sum_{\nu} n_\nu = O(n \log \frac{1}{\gamma})$. \square

We now state and prove a somewhat weaker analogue of the above lemma for the curved case.

Lemma 3.9. *Let $\tilde{U}(m)$ be the superadditive function measuring the maximum union complexity of any set of m vertical or m horizontal curved towers of constant complexity. Then the complexity of the union of any set \mathcal{F} of n locally γ -fat objects of constant complexity is $O(\frac{1}{\gamma^2} \cdot \tilde{U}(\frac{cn}{\gamma^2}))$, for some constant $c > 0$ that depends only on the maximum description complexity of individual objects; the overall constant of proportionality also depends on individual object complexity.*

Proof. The proof is similar to the proof of the previous lemma, the main difference being that we use the Merging Lemma instead of the Combination Lemma.

Let T be a set of n well-adjusted quasi-triangles. As in the above proof, we cover each well-adjusted quasi-triangle by three horizontal towers, three vertical towers, and six axis-aligned rectangles of aspect ratio bounded by a constant, thus obtaining sets T_{hor} , T_{vert} , and A of these three respective types of objects. Using the fact that the towers and rectangles have constant (local)

fatness, the Merging Lemma and the superadditivity of $\tilde{U}(\cdot)$ imply that the union complexity of T is bounded by

$$|\mathcal{U}(T)| \leq O(1) \cdot (|\mathcal{U}(T_{\text{hor}})| + |\mathcal{U}(T_{\text{vert}})|) + O(n) \leq O(\tilde{U}(7n)).$$

(The last inequality is due to the fact that, since $\tilde{U}(\cdot)$ measures the maximum union complexity of any set of curved towers, it must grow faster than any linear function. In other words, the maximum union complexity $\tilde{U}_w(n)$ of a set of n well-adjusted quasi-triangles satisfies $\tilde{U}_w(n) = O(\tilde{U}(7n))$.)

Now let \mathcal{F} be an arbitrary set of n locally γ -fat objects of constant complexity. By Lemma 3.4 we know that $|\mathcal{U}(\mathcal{F})| \leq |\mathcal{U}(\mathcal{G}_1 \cup \dots \cup \mathcal{G}_\kappa)|$, where each \mathcal{G}_i is a collection of right γ^* -quasi-triangles, with $\kappa = O(1/\gamma)$, $\gamma^* = \Theta(\gamma)$, and $\sum_i |\mathcal{G}_i| = O(n/\gamma^2)$. Moreover, Lemma 3.4(ii) implies that each collection \mathcal{G}_i can be transformed to a collection of well-adjusted quasi-triangles, so the union complexity of \mathcal{G}_i is at most $\tilde{U}_w(n_i) \leq O(\tilde{U}(7n_i))$, where $n_i := |\mathcal{G}_i|$. Note that each vertex of $\mathcal{U}(\mathcal{F})$ is a vertex of some $\mathcal{U}(\mathcal{G}_i \cup \mathcal{G}_j)$, and that $|\mathcal{U}(\mathcal{G}_i \cup \mathcal{G}_j)| = O(1/\gamma) \cdot (|\mathcal{U}(\mathcal{G}_i)| + |\mathcal{U}(\mathcal{G}_j)|)$ by the Merging Lemma and Lemma 3.3. Hence,

$$\begin{aligned} |\mathcal{U}(\mathcal{F})| &\leq \sum_{1 \leq i < \kappa} \sum_{i < j \leq \kappa} |\mathcal{U}(\mathcal{G}_i \cup \mathcal{G}_j)| \\ &= \sum_{1 \leq i < \kappa} \sum_{i < j \leq \kappa} O\left(\frac{1}{\gamma}\right) \cdot (|\mathcal{U}(\mathcal{G}_i)| + |\mathcal{U}(\mathcal{G}_j)|) \\ &\leq \sum_{1 \leq i < \kappa} \sum_{i < j \leq \kappa} O\left(\frac{1}{\gamma}\right) \cdot (\tilde{U}(7n_i) + \tilde{U}(7n_j)) \\ &\leq O\left(\frac{1}{\gamma}\right) \cdot \left(\sum_{1 \leq i < \kappa} \sum_{1 < j \leq \kappa} \tilde{U}(7n_i) + \sum_{1 \leq j < \kappa} \sum_{1 < i \leq \kappa} \tilde{U}(7n_j) \right) \\ &\leq O\left(\frac{1}{\gamma^2}\right) \cdot \left(\sum_{1 \leq i < \kappa} \tilde{U}(7n_i) + \sum_{1 \leq j < \kappa} \tilde{U}(7n_j) \right) && \text{(because } \kappa = O(1/\gamma)\text{)} \\ &\leq O\left(\frac{1}{\gamma^2}\right) \cdot \tilde{U}\left(14 \sum_{1 \leq i < \kappa} n_i\right) && \text{(because } \tilde{U}(\cdot) \text{ is superadditive)} \\ &\leq O\left(\frac{1}{\gamma^2}\right) \cdot \tilde{U}\left(\frac{cn}{\gamma^2}\right) \quad \text{for a suitable constant } c && \text{(because } \sum_i n_i = O(n/\gamma^2)\text{)}. \end{aligned}$$

□

4 The union complexity of towers

In this section we establish an upper bound on the complexity of the union of n vertical or n horizontal towers. Since the two cases are symmetric, we only handle the case of vertical towers. Thus, in what follows a “tower” always means a vertical tower.

Let \mathcal{T} be a set of n towers. Our goal is to bound the union complexity of \mathcal{T} . Up to an additive $O(n)$ term, this complexity is equal to the number of union vertices that arise as the intersections of edges of pairs of towers. We distinguish four types of such vertices: (i) An *HV-vertex* is the intersection of a horizontal edge and a vertical edge. (ii) An *HD-vertex* is the intersection of a horizontal edge and a diagonal edge. (iii) A *VD-vertex* is the intersection of a vertical edge and a diagonal edge. (iv) A *DD-vertex* is the intersection of two diagonal edges.

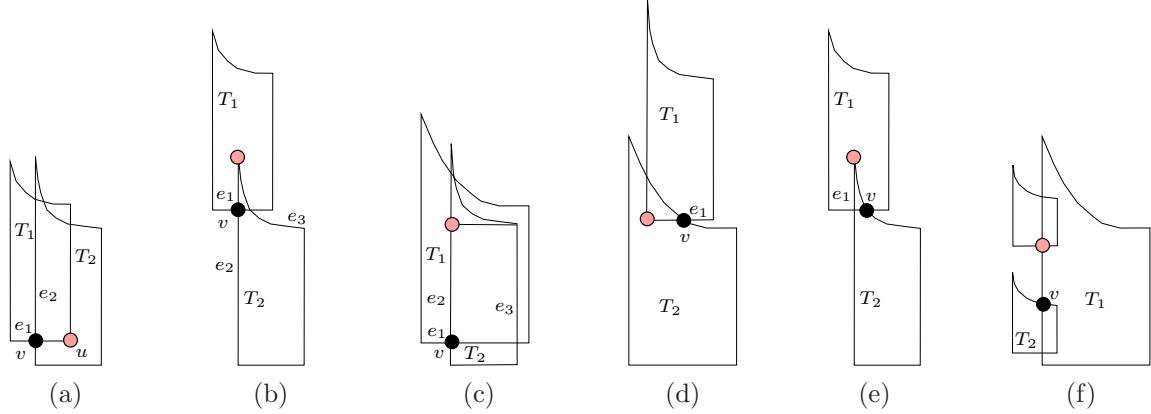


Figure 7: Charges made by HV-, HD-, and VD-vertices.

4.1 The number of HV-, HD-, and VD-vertices

We start by proving an $O(n)$ bound on the number of HV-, HD-, and VD-vertices. We will prove this bound for the case of curved towers, which immediately implies the same bound for the straight case.

Lemma 4.1. *Let \mathcal{T} be a set of n towers. Then the number of HV-, HD-, and VD-vertices of $\mathcal{U}(\mathcal{T})$ is $O(n)$.*

Proof. We first bound the number of HV-vertices. To this end we will charge each HV-vertex to a corner of some tower or pillar, in such a way that each corner gets charged at most $O(1)$ times; the linear bound will then follow immediately. The charging is done as follows.

Let $v = e_1 \cap e_2$ be an HV-vertex, where e_1 is the horizontal edge—that is, the bottom edge—of some tower T_1 and e_2 is a vertical edge of some tower T_2 . Refer to Figures 7(a)–(c). Starting at v , follow the edge e_1 into the interior of T_2 . We have two cases.

Case (i): e_1 ends before it exits T_2 . In this case one of the endpoints of e_1 lies inside T_2 . We charge v to this endpoint. An endpoint u of any bottom edge e gets charged at most once in this manner, since u can only be charged by the union vertex closest to u on e_1 . See Figure 7(a).

Case (ii): e_1 exits T_2 through some edge e_3 . Here we distinguish two subcases.

If e_3 is the diagonal of T_2 , then the top corner of T_2 must be inside T_1 . To see this, follow e_3 into T_1 from its intersection with the bottom edge of T_1 . The diagonal edge cannot exit T_1 through the left edge of T_1 (because the left edge of T_2 intersects the bottom edge of T_1), nor can it exit T_1 through its top edge (by Observation 3.5). We now charge v to the top corner of T_2 . Arguing as in Case (i), any top corner gets charged at most once in this manner. See Figure 7(b).

Now suppose both e_2 and e_3 are vertical edges. Then the two top corners of $\text{pillar}(T_2)$ are contained in $\text{pillar}(T_1)$, because $\text{pillar}(T_1)$ and $\text{pillar}(T_2)$ have the same aspect ratio (namely, 3). In particular, if we move from v into the interior of T_1 along e_2 we reach a top corner of

pillar(T_2) before we exit T_1 . We charge v to this pillar corner. Any pillar corner is charged at most once in this manner. See Figure 7(c).

Hence, the number of HV-vertices of the union is $O(n)$.

The analysis of HD-vertices proceeds similarly. If the horizontal edge e_1 incident to an HD-vertex v ends inside the other tower, we charge v to the corresponding endpoint of e_1 , as in Case (i) (see Figure 7(d)). Otherwise, we are in the same situation as in Case (ii) above, and then v charges the top endpoint of the incident diagonal edge (see Figure 7(e)). Thus any tower corner is charged at most once in this manner, and so the number of HD-vertices is $O(n)$.

It remains to bound the number of VD-vertices. Note that any VD-vertex is the lower endpoint of a vertical union edge e ; see Figure 7(f). The upper endpoint of e must either be an HV-vertex or the corner of the tower containing that edge on its boundary. Since the number of HV-vertices and the number of tower corners is $O(n)$, this implies that the number of VD-vertices is also $O(n)$. \square

It remains to bound the number of DD-vertices of $\mathcal{U}(\mathcal{T})$. To this end we first consider in Section 4.2 a special case, namely when all the towers in \mathcal{T} are stabbed by a common vertical line. This will be used in the analysis of the general case, which is given in Section 4.3.

4.2 The union complexity of vertical towers stabbed by a vertical line

The following lemma is the main result of this subsection.

Lemma 4.2. *Let \mathcal{T} be a set of n vertical towers that are all intersected by a vertical line ℓ . Then the complexity of $\mathcal{U}(\mathcal{T})$ is $O(\lambda_{s+2}(n))$, where s is the maximum number of times any pair of tower diagonals intersect. When $s = 1$ (which includes the straight case) this bound improves to $O(n)$.*

Note that any other pair of tower edges—that is, any pair with at least one non-diagonal edge—intersect at most once.

Proof. We first prove the bound $O(\lambda_{s+2}(n))$ for any $s \geq 1$, and then show the improvement for the case $s = 1$.

By Lemma 4.1, it suffices to prove that the number of DD-vertices of $\mathcal{U}(\mathcal{T})$ is $O(\lambda_{s+2}(n))$. Remove the pillars from the towers, that is, consider the set $\text{tops}(\mathcal{T}) = \{\text{top}(T) \mid T \in \mathcal{T}\}$. Observe that any DD-vertex of $\mathcal{U}(\mathcal{T})$ is also a DD-vertex of $\mathcal{U}(\text{tops}(\mathcal{T}))$.

The union $\mathcal{U}(\text{tops}(\mathcal{T}))$ may have holes. However, we will argue that any such hole must be completely covered by the pillars that we have removed. Thus, the only vertices of $\mathcal{U}(\text{tops}(\mathcal{T}))$ that can also appear on $\partial\mathcal{U}(\mathcal{T})$ are vertices on the outer face of the arrangement $\mathcal{A}(\text{tops}(\mathcal{T}))$. This outer face is contained in the outer face of the arrangement of the n diagonals of the tops in $\text{tops}(\mathcal{T})$, and every DD-vertex of the former face is also a vertex of the latter. In other words, we want to bound the complexity of a single cell in an arrangement of n curves, each intersecting the line ℓ , so that any pair of curves intersect in at most s points. The complexity of such a cell is $O(\lambda_{s+2}(n))$, as follows from the theory in [34, Theorem 5.3].

It remains to prove the claim that all holes of $\mathcal{U}(\text{tops}(\mathcal{T}))$ are covered by the removed pillars. First, note that there cannot be any portion of a hole to the right of ℓ , because the boundary of the union in this halfplane is weakly y -monotone; put differently, for each point on the union boundary, the horizontal segment connecting it to ℓ is fully contained in the union. Now consider a hole H

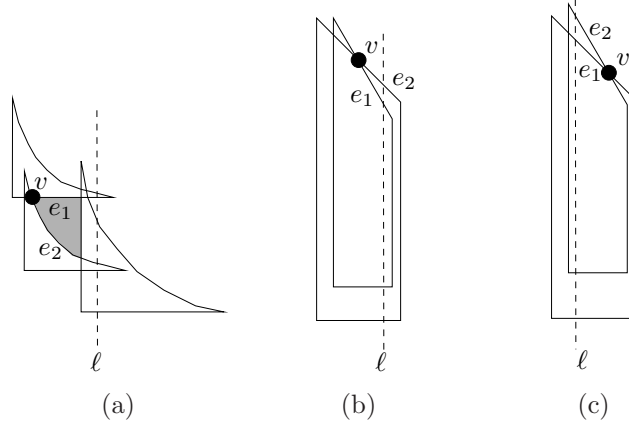


Figure 8: The proof of Lemma 4.3.

to the left of ℓ . Its leftmost vertex v must be the intersection of a horizontal edge e_1 of the top of some tower T_1 and the diagonal edge e_2 of the top of some tower T_2 . Since e_1 and e_2 both intersect ℓ , the hole H is contained in the region enclosed by e_1 , e_2 , and ℓ —see Figure 8(a). However, this region is completely covered by $\text{pillar}(T_1)$, because by Observation 3.5 the edge e_2 cannot intersect both the top edge and the bottom edge of $\text{pillar}(T_1)$. This establishes the bound $O(\lambda_{s+2}(n))$.

We are left to analyze the case $s = 1$. As above, it suffices to prove that the number of DD-vertices of $\mathcal{U}(\mathcal{T})$ is $O(n)$. First consider a DD-vertex v to the left of ℓ . Let e_1 and e_2 be the diagonal edges defining v , and let T_1, T_2 be their corresponding towers. Without loss of generality, suppose that e_1 lies below e_2 locally to the right of v . Since e_2 must intersect ℓ , this implies that v is the rightmost union vertex along e_1 to the left of ℓ . Indeed, this follows from Observation 3.5, as e_1 cannot meet the bottom edge of T_2 . Thus we charge v to e_1 , and any diagonal edge is charged only once in this manner. See Figure 8(b).

Now consider a DD-vertex v to the right of ℓ , incident to two diagonals e_1, e_2 , so that e_1 lies below e_2 locally to the left of v . Since e_2 intersects ℓ , this implies that v is the leftmost union vertex along e_1 to the right of ℓ —the portion of e_1 between ℓ and v is contained in the tower of e_2 . See Figure 8(c). This completes the proof of the lemma. \square

4.3 The global proof: recursive counting

We now return to the general case, where the towers in \mathcal{T} are not necessarily intersected by a common vertical line. Our strategy is to decompose the plane into cells in a suitable manner, and then recursively count the number of union vertices in each of the cells. This leads to a recurrence on the union complexity, whose solution yields the bounds mentioned in the introduction.

4.3.1 The decomposition

To obtain our decomposition, we first partition the set \mathcal{T} of towers into $O(\log n)$ subsets such that the union complexity of each subset is easy to bound, and then we construct the decomposition based on the union boundaries of these subsets.

The partitioning into subsets is done using a standard interval-tree approach [21, 31]. The root of the interval tree corresponds to a vertical line ℓ that splits \mathcal{T} into three subsets: a subset \mathcal{T}_L of

towers lying fully to the left of ℓ , a subset \mathcal{T}_R of towers lying fully to the right of ℓ , and a subset \mathcal{T}_\times of towers intersecting ℓ . The line ℓ is chosen such that $|\mathcal{T}_L| \leq |\mathcal{T}|/2$ and $|\mathcal{T}_R| \leq |\mathcal{T}|/2$. We store \mathcal{T}_\times at the root and continue to process \mathcal{T}_L at the left child and \mathcal{T}_R at the right child, recursively. The resulting tree has depth $k \leq \log n + 1$ and induces a partitioning of \mathcal{T} into subsets $\mathcal{T}_1, \mathcal{T}_2, \dots, \mathcal{T}_k$, where \mathcal{T}_j is the union of the subsets stored at nodes of level j . Now consider one such subset \mathcal{T}_j . It consists of subsets, one for each node at level j , such that the towers within each subset are stabbed by a common vertical line and any two towers from different subsets are disjoint. (In fact, their x -projections are disjoint). The following lemma thus follows directly from Lemma 4.2.

Lemma 4.3. *The complexity of the union $\mathcal{U}_j := \mathcal{U}(\mathcal{T}_j)$ of each subfamily \mathcal{T}_j is $O(\lambda_{s+2}(n_j))$, where s is the maximum number of times any pair of tower diagonals intersect and $n_j := |\mathcal{T}_j|$. When $s = 1$ (which includes the straight case) this bound improves to $O(n_j)$.*

We now obtain our decomposition as follows. Define E_j to be the set of edges of the union \mathcal{U}_j , for $j = 1, \dots, k$, and put $E := \bigcup_{i=1}^k E_i$. We next exclude from E (resp., from each E_j) all vertical edges, and with a slight abuse of notation continue to denote this set by E (resp., E_j) and the original set by \bar{E} (resp., \bar{E}_j , $j = 1, \dots, k$). Put $N := |E|$, and note that $N = O(\lambda_{s+2}(n))$ for the curved case, and $N = O(n)$ for the straight case. Construct a $(1/r)$ -cutting Ξ for E , that is, a partitioning of the plane into cells such that each cell σ is crossed by at most $|E|/r$ edges from E [13]. See Lemma 4.4 below for more details. Here r is a suitably chosen parameter—see below. The cells of the cutting are bounded from above and below by (portions of) edges from E , and from the left and the right by vertical segments. With a slight abuse of terminology (in the curved case) we will call such cells *trapezoids*. (Curved cells of this kind are generally called *pseudo-trapezoids*, but we prefer this shorthand notation for convenience.) Note that the left or right edge of a trapezoid may degenerate into a point.

Before we proceed to use the cutting to derive our recurrence, we first need to specify the value of r and bound the number of cells in the cutting. From now on we assume that each tower contributes at least one edge to the union boundary—towers for which this is not the case can simply be removed without decreasing the union complexity. Hence, we may assume that $N \geq n$. To simplify the notation, we define $\beta_{s+2}(m) := O(\lambda_{s+2}(m)/m)$, and set $\beta(s, m)$ to be $\beta_{s+2}(m)$ if $s > 1$, and $O(1)$ otherwise, choosing the constants of proportionality so that

$$|\mathcal{U}_j| = |\mathcal{U}(\mathcal{T}_j)| \leq n_j \beta(s, n_j) \leq n_j \beta(s, n),$$

for each family \mathcal{T}_j . Note that $N = O(\lambda_{s+2}(n)) = O(n\beta(s, n))$ for $s > 1$, and that $N = O(n) = O(n\beta(s, n))$ for $s = 1$ too. That is, we always have $N = O(n\beta(s, n))$. Note also that $\beta(s, n)$ is a very slowly growing function of n for any fixed s .

Lemma 4.4. *Let $r := \frac{N}{\beta(s, n) \log n}$. Then there exists a $(1/r)$ -cutting for E that consists of $O(r)$ trapezoidal cells and such that the top and bottom edges of each cell are contained in edges of E .*

Proof. Let X denote the total number of intersections among the edges in E . We claim (and will prove below) that $X = O(n\beta(s, n) \log n) = O(N\beta(s, n) \log n)$. The lemma then follows immediately from Theorem 2.1 of De Berg and Schwarzkopf [8], which asserts that there exists a $(1/r)$ -cutting for $\mathcal{A}(E)$ consisting of $O(r + Xr^2/N^2)$ cells. (Theorem 2.1 in [8] stating this bound is given for the case of straight line segments, but the proof also applies to curves. Moreover, it is easy to check

that the trapezoids in the cutting generated in [8] have the required property, namely that their top and bottom edges are contained in edges from the given set E .)

It remains to prove the bound on X . Recall that \bar{E} is E together with the vertical edges that we excluded. We in fact establish the same asymptotic upper bound on the number \bar{X} of intersections among the edges in \bar{E} . Consider an intersection point v of two edges $e, e' \in \bar{E}$, and assume v is not an endpoint of e or e' . Let j, j' be such that $e \in \bar{E}_j$ and $e' \in \bar{E}_{j'}$. In other words, e is a boundary edge of the union $\mathcal{U}(\mathcal{T}_j)$ and e' is a boundary edge of the union $\mathcal{U}(\mathcal{T}_{j'})$. We have $j \neq j'$, because the relative interiors of edges of the same union are pairwise disjoint by definition. Moreover, v must be a vertex of $\mathcal{U}(\mathcal{T}_j \cup \mathcal{T}_{j'})$. Using the Merging Lemma and the fact that towers have constant local fatness (see Lemma 3.6) together with Lemma 4.3 and the superadditivity of λ_{s+2} , we conclude that the complexity of $\mathcal{U}(\mathcal{T}_j \cup \mathcal{T}_{j'})$ is at most $O(\lambda_{s+2}(n_j + n_{j'}))$ for the curved case, and $O(n_j + n_{j'})$ for the straight case, which is $O((n_j + n_{j'})\beta(s, n_j + n_{j'}))$ in both cases. Summing over all pairs of unions $\mathcal{U}(\mathcal{T}_j), \mathcal{U}(\mathcal{T}_{j'})$ such that $j < j'$, and noting that the number of unions is only $k = O(\log n)$, we conclude that

$$\begin{aligned} \bar{X} &= \sum_{1 \leq j < j' \leq k} \sum O((n_j + n_{j'})\beta(s, n_j + n_{j'})) \\ &= O\left(\sum_{1 \leq j \leq k} (n_j k + n)\beta(s, n)\right) = O(nk\beta(s, n)) \\ &= O(n\beta(s, n) \log n). \end{aligned} \quad \square$$

Remark. In view of the bound $N = O(n\beta(s, n))$, we have $r = O(n/\log n)$.

4.3.2 Setting up the recurrence

Our goal is to derive a recurrence on the maximum union complexity of any set of n vertical towers, based on the cutting provided by Lemma 4.4 and the results from the previous sections. (Obviously, the resulting bound will hold for n horizontal towers as well.) To get good bounds from the recurrence, we need to keep the total number of towers in the recursion under control. Specifically, the towers that participate in a subproblem associated with a trapezoid σ of the cutting are classified according to the way in which they intersect σ . As it turns out (and spelled out in detail below), the only class of towers whose union needs further recursive processing are those that have a vertex inside σ ; the other towers that cross σ will be handled at this step and will not be passed on further down the recursion. Thus a tower T may be sent to up to four recursive subproblems, but in fact we have the stronger property that at any recursion depth the overall number of subproblems that involve T in this manner is at most four. To handle this bookkeeping properly, we introduce the notion of a *weight* of a tower T , denoted as $\text{weight}(T)$, which is equal to the number of vertices of T inside the trapezoid associated with the current subproblem; so $\text{weight}(T)$ is an integer between 1 and 4. The weight of the subproblem at σ is the sum of the weights of its towers. It follows that at each recursive step the weight is partitioned, without duplications, into the various subproblems. In other words, the total weight of all subproblems at any fixed recursion level never exceeds $4n$, where n is number of towers in the original problem, and this lets us keep the size of the recursive problems under control.

To exploit this observation, we define the *size* of a subproblem to be the total weight w of its towers rather than their number m ; note that we always have $m \leq w \leq 4m$ (recall that m counts

only the towers that have at least one vertex in the trapezoid σ). The recurrence that we develop has a non-recursive “overhead” term that depends on m but, replacing m by w , we transform the overhead term to (a potentially larger) term depending on w . The weight is then partitioned among the recursive subproblems, and we can therefore express the right-hand side of the recurrence in terms of the weights alone.

This suggests that instead of bounding the maximum complexity $U(n)$ of the union of n (vertical) towers, we bound the maximum complexity $U_\omega(w)$ of vertical towers with total weight w , and then apply the inequality $U(n) \leq U_\omega(4n)$ to obtain the bound that we really want.

Now consider a trapezoid σ and a set \mathcal{T} of n (straight or curved) towers with total weight w . As just discussed, the only parameter that we want to use in the recurrence that we are about to develop is w , but for the time being we ignore w completely, and treat the subproblem as one involving $n \leq w$ (unweighted) towers. The parameter w will re-enter the analysis only towards the very end.

We first apply the decomposition described in Section 4.3.1 to \mathcal{T} ; that is, we construct a $(1/r)$ -cutting as described above: we partition \mathcal{T} into $k = \log n + 1$ sets $\mathcal{T}_1, \dots, \mathcal{T}_k$ using an interval-tree approach, and then construct a $(1/r)$ -cutting Ξ for the set $E := \bigcup_{i=1}^k E_i$, where E_i is the set of (non-vertical) edges of $\mathcal{U}(\mathcal{T}_i)$. Here we clip the edges of E to within the trapezoid σ , so that we can ensure that the cutting Ξ is a subdivision of σ . We set the parameter r , as in Lemma 4.4, to $N/(\beta(s, n) \log n)$, where $N := |E|$. Recall that $N = O(n\beta(s, n))$, so, as already noted, $r = O(n/\log n)$. Moreover, $N/r = O(\beta(s, n) \log n)$, which is $O(\log n)$ for $s = 1$.

In what follows we drop the notation $\beta(s, n)$ and spell out the different bounds for $s > 1$ and for $s = 1$. According to Lemma 4.4 our cutting Ξ thus has the following properties:

- In the straight case, Ξ consists of $O(n/\log n)$ trapezoidal cells, each intersected by at most $O(\log n)$ edges of E .
- In the curved case, Ξ consists of $O(n/\log n)$ trapezoidal cells, each intersected by at most $O(\beta_{s+2}(n) \log n)$ edges of E .

Let $\sigma_1, \dots, \sigma_{|\Xi|}$ be the cells of the cutting Ξ . Each union vertex will lie in one of the cells—in fact, in at least one, since a vertex can lie on a shared boundary of two or more cells—so it suffices to recursively count the union complexity inside each cell σ_i , including the vertices on its boundary, and sum the results. (Alternatively, we can dispose of the vertices on the boundary immediately, by observing that their number is bounded by $O(|\Xi| \cdot N/r) = O(N)$.) Within a cell σ_i we only recurse on the towers that have a vertex in the interior of σ_i ; union vertices induced by other towers will be counted separately, and added as an overhead term in the recursion. Thus at each recursive step we construct a subdivision within σ_i according only to the towers that have a vertex in its interior. We clip the resulting trapezoids to σ_i , and, when necessary, further refine them to become actual trapezoids (confined to σ_i). Next we make this idea precise.

Consider a cell σ_i of Ξ . We say that a tower T is *visible* within σ_i if at least one edge of E lying on ∂T crosses σ_i . Otherwise, T is *hidden* within σ_i , in the sense that its boundary crosses σ but only within the interior of the union. (Here and hereafter, a curve “crossing” a region refers to the curve intersecting the interior of the region.) Since any union vertex of type DD lies on at least one edge from E —more precisely, any DD-vertex of $\mathcal{U}(\mathcal{T})$ is either an endpoint of an edge in E or an intersection of two such edges—we only have to consider visible towers within σ_i . We note that a tower T hidden in σ_i may still have a vertical edge crossing σ_i and appearing on the

boundary of the union. In this case we can still discard T , as it does not contribute a DD-vertex to the union, and the total number of union vertices that lie on such edges is only linear (as shown in Lemma 4.1). Discarding invisible towers can expose vertices that were not on the union boundary before. This, however, only leads to overcounting in the recursion, which is consistent with our aim for an upper bound on the union complexity. Let $\mathcal{T}(\sigma_i)$ denote the set of towers that are visible within σ_i , and set $n_i := |\mathcal{T}(\sigma_i)|$. Since $|\mathcal{T}(\sigma_i)|$ is at most the number of edges of E that cross σ_i , we have $n_i \leq N/r$ and $\sum_i n_i = O(r \cdot (N/r)) = O(N)$.

We further partition $\mathcal{T}(\sigma_i)$ into four (not necessarily disjoint) subsets:

- a subset $\mathcal{T}_{\text{vert}}(\sigma_i)$ containing all towers $T \in \mathcal{T}(\sigma_i)$ that have a vertex in the interior of σ_i ;
- a subset $\mathcal{T}_{\text{left}}(\sigma_i)$ of towers intersecting the line supporting the left side of σ_i ;
- a subset $\mathcal{T}_{\text{right}}(\sigma_i)$ of towers intersecting the line supporting the right side of σ_i ;
- a subset $\mathcal{T}_{\text{rest}}(\sigma_i) = \mathcal{T}(\sigma_i) \setminus (\mathcal{T}_{\text{vert}}(\sigma_i) \cup \mathcal{T}_{\text{left}}(\sigma_i) \cup \mathcal{T}_{\text{right}}(\sigma_i))$ containing the remaining towers.

Note that $\sum_{1 \leq i \leq |\Xi|} \text{weight}(\mathcal{T}_{\text{vert}}(\sigma_i)) \leq w$; this crucial property will be used towards the end of the analysis, but we ignore it in the present part, in which tower weights play no role. Note also that, by definition, the towers in $\mathcal{T}_{\text{left}}(\sigma_i)$ (resp., $\mathcal{T}_{\text{right}}(\sigma_i)$) do not necessarily meet the left (resp., right) side of σ_i , but only the line supporting it. We first bound the union complexities of the individual sets and then consider their interaction.

A bound on the union complexities of $\mathcal{T}_{\text{left}}(\sigma_i)$ and $\mathcal{T}_{\text{right}}(\sigma_i)$ directly follows from Lemma 4.2.

Lemma 4.5. *We have $|\mathcal{U}(\mathcal{T}_{\text{left}}(\sigma_i))| = O(\lambda_{s+2}(n_i))$ and $|\mathcal{U}(\mathcal{T}_{\text{right}}(\sigma_i))| = O(\lambda_{s+2}(n_i))$, and in the straight case $|\mathcal{U}(\mathcal{T}_{\text{left}}(\sigma_i))| = O(n_i)$ and $|\mathcal{U}(\mathcal{T}_{\text{right}}(\sigma_i))| = O(n_i)$.*

The next lemma bounds the complexity of $\mathcal{U}(\mathcal{T}_{\text{rest}}(\sigma_i))$. In particular, note that the analysis for the straight case is somewhat different from that of the curved case, and thus cannot be derived directly from that case.

Lemma 4.6. *For the straight case we have $|\mathcal{U}(\mathcal{T}_{\text{rest}}(\sigma_i))| = O(n_i)$, and for the curved case we have $|\mathcal{U}(\mathcal{T}_{\text{rest}}(\sigma_i))| = O(\lambda_{s+2}(n_i))$.*

Proof. Any tower $T \in \mathcal{T}_{\text{rest}}(\sigma_i)$ must lie in the vertical strip delimited by the lines supporting the left and right sides of σ_i . We partition the set $\mathcal{T}_{\text{rest}}(\sigma_i)$ into two subsets, defined momentarily and denoted $\mathcal{T}_{\text{rest}}^{(1)}(\sigma_i)$ and $\mathcal{T}_{\text{rest}}^{(2)}(\sigma_i)$. We first bound their union complexities (inside σ_i) separately, and then analyze the complexity of the merged union. For our analysis it is crucial that the top and bottom edges of the trapezoid σ_i are contained in edges of E and, hence, in a diagonal or bottom edge of some tower. This implies that these edges are either horizontal or xy -monotone. Moreover, such edges cannot cross both the top edge and the bottom edge of the pillar of any tower—see Observation 3.5.

The subset $\mathcal{T}_{\text{rest}}^{(1)}(\sigma_i)$ consists of the towers $T \in \mathcal{T}_{\text{rest}}(\sigma_i)$ whose diagonal edge $e_{\text{dia}}(T)$ lies entirely above the interior of σ_i .

If $e_{\text{dia}}(T)$ lies entirely above σ_i then $T \cap \sigma_i = B(T) \cap \sigma_i$, where $B(T)$ is the axis-aligned bounding box of T . Hence, for the purpose of bounding the union complexity inside σ_i , we can replace T by $B(T)$; see Figure 9(a). Note that $B(T)$ has constant (albeit not fixed) aspect ratio because T is locally $\Omega(1)$ -fat. Since the union of a set of constant aspect-ratio axis-parallel rectangles has linear

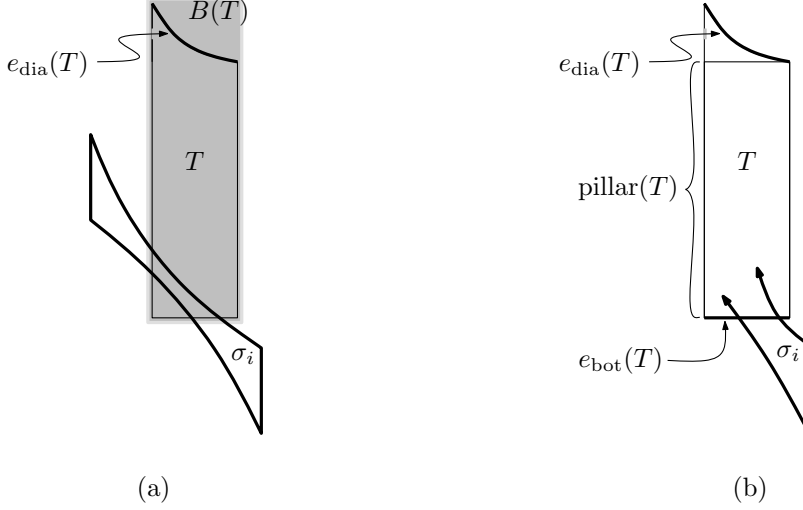


Figure 9: (a) The diagonal edge $e_{\text{dia}}(T)$ lies above σ_i , so we can replace T by its bounding box $B(T)$. (b) The top edge of σ_i cannot exit $\text{pillar}(T)$ through the top of $\text{pillar}(T)$.

complexity, this implies that the complexity of the union of $\mathcal{T}_{\text{rest}}^{(1)}(\sigma_i)$ inside σ_i is $O(|\mathcal{T}_{\text{rest}}^{(1)}(\sigma_i)|) = O(n_i)$.

The subset $\mathcal{T}_{\text{rest}}^{(2)}(\sigma_i)$ consists of the remaining towers from $\mathcal{T}_{\text{rest}}(\sigma_i)$. Note that the diagonal edge of any such tower must intersect σ_i : if it lay completely below σ_i then T would lie completely below σ_i , contradicting the fact that T is visible in σ_i .

Claim 4.7. *The bottom edge $e_{\text{bot}}(T)$ of any tower $T \in \mathcal{T}_{\text{rest}}^{(2)}(\sigma_i)$ lies completely below σ_i .*

Proof. To prove the claim we will show that both endpoints of $e_{\text{bot}}(T)$ lie below σ_i . Since the bottom edge of σ_i is horizontal or xy -monotone, it cannot cross $e_{\text{bot}}(T)$ twice. Thus $e_{\text{bot}}(T)$ lies completely below σ_i if both endpoints lie below σ_i .

Observe that at least one endpoint of $e_{\text{bot}}(T)$ must lie below σ_i : by definition, neither of them lies inside σ_i , and if both endpoints lay above σ_i then T would lie completely above σ_i as well, since the top edge of σ_i cannot cross $e_{\text{bot}}(T)$ twice. Now suppose for a contradiction that $e_{\text{bot}}(T)$ also has an endpoint above σ_i . Then the top edge of σ_i must cross $e_{\text{bot}}(T)$, as in Figure 9(b). By Observation 3.5 this edge cannot also cross the top edge of $\text{pillar}(T)$ —it must exit $\text{pillar}(T)$ through its left side. But then the diagonal edge of T lies completely above σ_i , contradicting the fact that $T \in \mathcal{T}_{\text{rest}}^{(2)}(\sigma_i)$. \square

In the straight case $e_{\text{dia}}(T)$ must have exactly one vertex, v , above σ_i , since having both vertices above (resp., below) σ_i would imply in the straight case that $e_{\text{dia}}(T)$ lies completely above (resp., below) σ_i . For the straight case Claim 4.7 now implies that $T \cap \sigma_i = W(T) \cap \sigma_i$, where $W(T)$ is the wedge that has v as an apex and is bounded by the rays emanating from v and containing the edges of T incident to v ; refer to Figure 10(a). Note that $W(T)$ is an $\Omega(1)$ -fat wedge because T is locally $\Omega(1)$ -fat. Thus the complexity of the union of $\mathcal{T}_{\text{rest}}^{(2)}(\sigma_i)$ inside σ_i is bounded by the complexity of the union of $|\mathcal{T}_{\text{rest}}^{(2)}(\sigma_i)| = O(n_i)$ fat wedges, which is $O(n_i)$ [18].

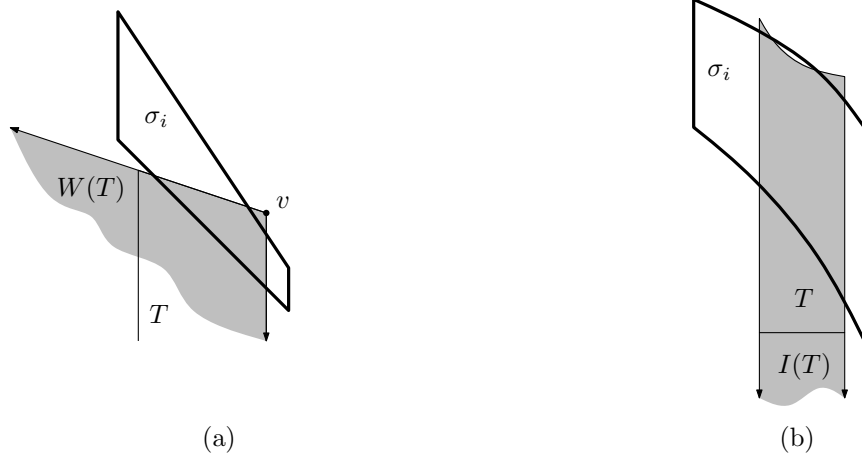


Figure 10: Situations in which the edge $e_{\text{bot}}(T)$ lies below σ_i : (a) In the straight case we can replace T by the wedge $W(T)$. (b) In the curved case we can replace T by the semi-infinite tower $I(T)$.

For the curved case Claim 4.7 implies that $T \cap \sigma_i = I(T) \cap \sigma_i$, where $I(T)$ is the “semi-infinite tower” obtained by extending the vertical edges of T downward to infinity and removing its bottom edge; refer to Figure 10(b). Thus the complexity of the union of $\mathcal{T}_{\text{rest}}^{(2)}(\sigma_i)$ inside σ_i is bounded by the complexity of the union of $|\mathcal{T}_{\text{rest}}^{(2)}(\sigma_i)| = O(n_i)$ such semi-infinite towers. This in turn is equal to the complexity of the upper envelope of the diagonal edges of the towers, which is $O(\lambda_{s+2}(n_i))$ [34].

We have bounded the individual union complexities of $\mathcal{T}_{\text{rest}}^{(1)}(\sigma_i)$ and $\mathcal{T}_{\text{rest}}^{(2)}(\sigma_i)$. The lemma now follows from the observation that $|\mathcal{U}(\mathcal{T}_{\text{rest}}^{(1)}(\sigma_i) \cup \mathcal{T}_{\text{rest}}^{(2)}(\sigma_i))| = O(|\mathcal{U}(\mathcal{T}_{\text{rest}}^{(1)}(\sigma_i))| + |\mathcal{U}(\mathcal{T}_{\text{rest}}^{(2)}(\sigma_i))|)$, which is an immediate consequence of the Merging Lemma (which can be applied in both the straight and curved cases). \square

It remains to consider the interaction between $\mathcal{T}_{\text{vert}}(\sigma_i)$, $\mathcal{T}_{\text{left}}(\sigma_i)$, $\mathcal{T}_{\text{right}}(\sigma_i)$, and $\mathcal{T}_{\text{rest}}(\sigma_i)$. First, consider $\mathcal{U}(\mathcal{T}_{\text{left}}(\sigma_i) \cup \mathcal{T}_{\text{right}}(\sigma_i) \cup \mathcal{T}_{\text{rest}}(\sigma_i))$. Any vertex of this union is a vertex of the union of a pair of these sets. As above, the Merging Lemma implies that the complexity of the union of any pair of these sets is linear in the sum of the individual complexities. We conclude that in the straight case

$$|\mathcal{U}(\mathcal{T}_{\text{left}}(\sigma_i) \cup \mathcal{T}_{\text{right}}(\sigma_i) \cup \mathcal{T}_{\text{rest}}(\sigma_i))| = O(n_i)$$

and in the curved case

$$|\mathcal{U}(\mathcal{T}_{\text{left}}(\sigma_i) \cup \mathcal{T}_{\text{right}}(\sigma_i) \cup \mathcal{T}_{\text{rest}}(\sigma_i))| = O(\lambda_{s+2}(n_i)).$$

Now consider $|\mathcal{U}(\mathcal{T}(\sigma_i))|$. We have

$$\mathcal{U}(\mathcal{T}(\sigma_i)) = \mathcal{U}(\mathcal{T}_{\text{vert}}(\sigma_i)) \cup \mathcal{U}(\mathcal{T}_{\text{left}}(\sigma_i) \cup \mathcal{T}_{\text{right}}(\sigma_i) \cup \mathcal{T}_{\text{rest}}(\sigma_i)).$$

In the curved case, another application of the Merging Lemma yields

$$|\mathcal{U}(\mathcal{T}(\sigma_i))| = O(|\mathcal{U}(\mathcal{T}_{\text{vert}}(\sigma_i))| + \lambda_{s+2}(n_i)).$$

In the straight case we can apply the Few Holes Lemma and the Combination Lemma. The Few Holes Lemma is stated for fat triangles, but we can also apply it here because a fat rectangle can be covered by two fat triangles and we can replace the fat wedges by sufficiently large fat triangles without decreasing the union complexity or the number of holes. We then obtain

$$|\mathcal{U}(\mathcal{T}(\sigma_i))| = |\mathcal{U}(\mathcal{T}_{\text{vert}}(\sigma_i))| + O(n_i).$$

This sharper form in the straight case (recall that in the curved case the first term was $O(|\mathcal{U}(\mathcal{T}_{\text{vert}}(\sigma_i))|)$ rather than exactly $|\mathcal{U}(\mathcal{T}_{\text{vert}}(\sigma_i))|$) is due to the stronger bound in the Combination Lemma as compared to that of the Merging Lemma.

We have bounded the complexity of $\mathcal{U}(\mathcal{T}(\sigma_i))$ within σ_i . The set $\mathcal{T}(\sigma_i)$ only contains towers that are visible in σ_i but, as we have already observed, this is sufficient to bound the union complexity.

Finally, we sum over all cells σ_i , and re-introduce the tower weights. That is, according to the strategy outlined at the beginning of the analysis, for each trapezoid σ_i , we upper bound $|\mathcal{U}(\mathcal{T}_{\text{vert}}(\sigma_i))|$ by $U_\omega(w_i)$, where $w_i = \text{weight}(\mathcal{T}_{\text{vert}}(\sigma_i))$. We then upper bound the overhead term. For this we recall that $n_i \leq N/r$ for each $i = 1, \dots, |\Xi|$, and $\sum_{1 \leq i \leq |\Xi|} n_i = O(N)$. Recall also that $N = O(n)$ for the straight case and $N = O(\lambda_{s+2}(n))$ for the curved case. It follows that the overhead term is $O(N) = O(n)$ in the straight case, and $O(\lambda_{s+2}(N)) = O(\lambda_{s+2}(\lambda_{s+2}(n))) = O(n\beta_{s+2}^2(n))$ in the curved case (the superadditivity of $\lambda_{s+2}(\cdot)$ is used in the latter bound). We then transform this overhead bound to a potentially larger bound by replacing n by w (recalling that $n \leq w$). Once the right-hand side of the recurrence is expressed in terms of w only, we can replace $|\mathcal{U}(\mathcal{T})|$ on the left-hand side by its maximum value $U_\omega(w)$, where w is the overall weight of the input towers, and, at long last, obtain our recurrences.

Lemma 4.8. *For the straight case we have*

$$U_\omega(w) \leq \left(\sum_{1 \leq i \leq |\Xi|} U_\omega(w_i) \right) + O(w),$$

where $\sum_{1 \leq i \leq |\Xi|} w_i \leq w$ and $w_i = O(\log w)$ for all i .

Lemma 4.9. *For the curved case we have*

$$\tilde{U}_\omega(w) = O \left(\sum_{1 \leq i \leq |\Xi|} \tilde{U}_\omega(w_i) \right) + O(w\beta_{s+2}^2(w)),$$

where $\sum_{1 \leq i \leq |\Xi|} w_i \leq w$ and $w_i = O(\beta_{s+2}(w) \log w)$ for all i .

4.3.3 Putting it all together

We are finally ready to prove our main results.

Theorem 4.10. *The complexity of the union of n γ -fat triangles in the plane is*

$$O \left(n \log^* n + \frac{n}{\gamma} \log^2 \frac{1}{\gamma} \right).$$

Proof. Consider the recurrence from Lemma 4.8. Since $\text{weight}(\mathcal{T}_{\text{vert}}(\sigma_i)) = O(\log w)$ for all i , the depth of the recurrence is $O(\log^* w)$. Moreover, the fact that $\sum_{1 \leq i \leq |\Xi|} \text{weight}(\mathcal{T}_{\text{vert}}(\sigma_i)) \leq w$ means that at each level of the recurrence the total weight of towers under consideration is bounded by w . Since the overhead term in the recurrence is linear and the depth of the recurrence is $O(\log^* w)$, this implies that $U_\omega(w) = O(w \log^* w)$. Because $U(n) \leq U_\omega(4n)$ we can now use Lemma 3.8 to finish the proof. \square

Remark: We do not know whether the dependence on γ in this bound is best possible. We note that Pach and Tardos [32] obtain a smaller factor involving γ , namely $\frac{1}{\gamma} \log \frac{1}{\gamma}$, but it multiplies a super-linear expression in n , whereas in our bound only the linear term depends on γ . At any rate, the best known lower bound on such a factor is only $\Omega(\frac{1}{\gamma})$, so the question remains open for both bounds, Pach and Tardos' and ours.

Theorem 4.11. *The complexity of the union of n locally γ -fat objects of constant complexity is at most $\frac{n}{\gamma^4} 2^{O(\log^* n)}$, where the constant of proportionality in the exponent $O(\log^* n)$ depends on the description complexity of the given objects, but not on γ .*

Proof. Consider the recurrence from Lemma 4.9. Let C be the constant of proportionality in the first term of the recurrence, so that

$$\tilde{U}_\omega(w) \leq C \cdot \left(\sum_{1 \leq i \leq |\Xi|} \tilde{U}_\omega(w_i) \right) + O(w \beta_{s+2}^2(w)).$$

Since the towers are well-adjusted, C does not depend on γ , although it does depend on s . Here $w_i = \text{weight}(\mathcal{T}_{\text{vert}}(\sigma_i)) = O(\beta_{s+2}(w) \log w)$ for all i , which implies that the depth of the recursion is $O(\log^* w)$ in this case as well. (This can be seen, e.g., by defining $f(w) = \beta_{s+2}(w) \log w$ and observing that $f(f(w)) = O(\log w)$.) We next unwind the recurrence, and stop as soon as w is smaller than some appropriately chosen positive absolute constant. The sum of the overhead terms at depth j is bounded by $O(C^j w \beta_{s+2}^2(w))$, so

$$\tilde{U}_\omega(w) = O(C^{O(\log^* w)} \cdot w \beta_{s+2}^2(w)) = w 2^{O(\log^* w)},$$

where the constant of proportionality in the exponent depends on C and thus on s ; it also depends on s for subsuming the factor $\beta_{s+2}^2(w)$. Using the fact that $\tilde{U}(n) \leq \tilde{U}_\omega(4n)$, we now plug this bound into Lemma 3.9 to conclude that the union complexity of n locally γ -fat objects of constant complexity is bounded by $\frac{n}{\gamma^4} 2^{O(\log^*(cn/\gamma^2))}$, where c is the constant appearing in that lemma. For $n > 1/\gamma$ we have $\log^*(cn/\gamma^2) < \log^*(cn^3) = O(\log^* n)$, and we get a bound of $\frac{n}{\gamma^4} 2^{O(\log^* n)}$. The same bound holds for $n \leq 1/\gamma$, because the union complexity is trivially bounded by $O(n^2)$. \square

5 Applications

Set cover and ε -nets. A *range space* (P, \mathcal{R}) is a pair consisting of an underlying universe P of objects, and a certain collection $\mathcal{R} \subseteq 2^P$ of its subsets (called *ranges*). We assume in this discussion that P and \mathcal{R} are finite. The *set cover* problem for (P, \mathcal{R}) is to find a minimum-size subcollection $S \subseteq \mathcal{R}$ whose union covers P . Put $\text{OPT} := |S|$. The general problem is known to be NP-hard and the problem remains NP-hard in most geometric settings [3]. The standard greedy algorithm [11] yields

a set cover with an approximation factor $O(\log |P|)$ without any further assumptions on the range space. This factor is asymptotically best possible to achieve in polynomial time, under appropriate complexity-theoretic assumptions [4, 23]. However, the approximation factor can be improved to $O(\log \text{OPT})$, still achievable in expected polynomial time, for many geometric scenarios [10, 12, 20].

This improvement is closely related to the notion of ε -nets in a “dual” variant of such a range space. Specifically, in this dual context, a subset $N \subseteq \mathcal{R}$ is called a *dual ε -net* for (P, \mathcal{R}) , for some given $0 < \varepsilon < 1$, if every point $p \in P$ which is contained in more than $\varepsilon|\mathcal{R}|$ ranges of \mathcal{R} is contained in a range of N . The classical result of Haussler and Welzl [26], specialized to this context, asserts that, if (P, \mathcal{R}) has so-called *finite VC-dimension* (which it does in most geometric contexts; see below), then it admits dual ε -nets of the above kind of size $O(\frac{1}{\varepsilon} \log \frac{1}{\varepsilon})$, where the constant of proportionality depends linearly on the VC-dimension.

If (P, \mathcal{R}) admits dual ε -nets of size $O(\frac{1}{\varepsilon} \varphi(\frac{1}{\varepsilon}))$, then one can obtain approximation factor $O(\varphi(\text{OPT}))$ in expected polynomial time for the corresponding set cover problem [10, 20]. Thus, since we always have $\varphi(\frac{1}{\varepsilon}) = O(\log \frac{1}{\varepsilon})$, the “default” result, for range spaces of finite VC-dimension, is an approximation factor $O(\log \text{OPT})$, as mentioned above.

As observed in several recent papers [3, 14, 36], the size of dual ε -nets is strongly related to the complexity of the union of the given ranges. Improving upon a slightly weaker bound given in [14], the following result was established by three of the authors [3] (see also [36]): Let P be a finite set of points in the plane, and let \mathcal{R} be a collection of planar regions of *constant complexity*, such that the complexity of the union of any subset of r of these regions is $O(r\phi(r))$. Then there exist dual ε -nets for (P, \mathcal{R}) whose size is $O(\frac{1}{\varepsilon} \log \phi(\frac{1}{\varepsilon}))$. This in turn yields the approximation factor (still computable in expected polynomial time) $O(\log \phi(\text{OPT}))$ for the corresponding set cover problem.

In the *weighted* set-cover problem each element in \mathcal{R} is assigned a non-negative weight, and the goal is to find a minimum-weight subcollection of \mathcal{R} whose union covers P . It has recently been shown by Varadarajan [37] that with the above property of the union, one can achieve an approximation factor to the minimum weight of $2^{O(\log^* n)} \log \phi(n)$ in polynomial time, where $n = |\mathcal{R}|$. Note that the strength of this bound is the fact that it depends only on the number of input objects and not on the optimal weight OPT , which in some cases might be significantly larger than n , in contrast to the unweighted case. Very recently, the approximation factor has been improved to $O(\max\{1, \log \phi(n)\})$ by Chan *et al.* [5], matching the bound for the unweighed case. Incorporating the bound of Theorems 4.10 and 4.11 we obtain the following three corollaries.

Theorem 5.1. *Any range space of points and locally γ -fat objects of constant complexity in the plane admits a dual ε -net of $O(\frac{1}{\varepsilon} \log^* \frac{1}{\varepsilon})$ objects, for any $0 < \varepsilon \leq 1$. When the objects are γ -fat triangles in the plane, the ε -net size improves to $O(\frac{1}{\varepsilon} \log \log^* \frac{1}{\varepsilon})$.*

Theorem 5.2. *There exists a randomized expected polynomial-time algorithm that, given a set P of points in the plane and a set \mathcal{T} of locally γ -fat objects of constant complexity that cover P , computes a set cover $\mathcal{T}' \subseteq \mathcal{T}$ for P of size $O(\text{OPT} \log^* \text{OPT})$, where OPT is the size of the smallest such cover. When \mathcal{T} is a set of γ -fat triangles in the plane, the set cover size improves to $O(\text{OPT} \log \log^* \text{OPT})$.*

Theorem 5.3. *There exists a randomized expected polynomial-time algorithm that, given a set P of points in the plane and a set \mathcal{T} of n locally γ -fat objects of constant complexity that cover P , each of which is assigned a non-negative weight, computes a set cover $\mathcal{T}' \subseteq \mathcal{T}$ for P of weight $O(\text{OPT} \log^* n)$, where OPT is the minimum weight of such a cover. When \mathcal{T} is a set of γ -fat triangles in the plane, the set cover weight improves to $O(\text{OPT} \log \log^* n)$.*

Hidden surface removal. In a typical hidden surface removal problem, we are given a set of n pairwise-disjoint objects in 3-space, and a viewing point v , and the goal is to construct the view of the given scene, as seen from v . This view consists of a subdivision of the viewing plane into maximal connected regions in each of which (a portion of) a single object, or no object can be seen. The resulting structure is the so-called *visibility map* of the given objects as seen from v .

Katz *et al.* [27] have proposed an output-sensitive algorithm to construct the visibility map, whose running time is $O((U(n) + k) \log^2 n)$, where $U(r)$ is an upper bound on the complexity of the union of the projections to the viewing plane of any r of the input objects (here $U(\cdot)$ is also assumed to be superadditive), and k is the complexity of the visibility map. The space requirement in this case is $O(U(n) \log n)$.

When the objects are horizontal locally γ -fat objects (resp., γ -fat triangles), and the viewing plane is also horizontal, the projections of the input objects (resp., triangles) onto the viewing plane are also γ -fat. (This also holds, with a smaller γ , when the objects are not too tilted with respect to the viewing plane.) Hence, applying the bounds in Theorems 4.10 and 4.11, we obtain the following improved bounds on the performance of the algorithm in [27] for our fat objects.

Theorem 5.4. *Given a set of n pairwise-disjoint horizontal locally γ -fat objects of constant complexity in 3-space, the view of this set from any fixed point v can be computed in $O((\frac{n}{\gamma^4} 2^{c \log^* n} + k) \log^2 n)$ time, using $O(\frac{n}{\gamma^4} 2^{c \log^* n} \log n)$ space, where k is the complexity of the visibility map and c is a constant depending only on the maximum complexity of the individual objects. The corresponding bounds improve to $O((n(\log^* n + \frac{1}{\gamma} \log^2 \frac{1}{\gamma}) + k) \log^2 n)$, $O(n(\log^* n + \frac{1}{\gamma} \log^2 \frac{1}{\gamma}) \log n)$, if the objects are γ -fat triangles.*

Conflict-free coloring. In *conflict-free coloring* (CF-coloring, for short) problems we are given a finite family \mathcal{R} of n regions of some fixed type (such as disks, pseudo-disks, axis-parallel rectangles, etc.), and the goal is to find the smallest integer k , such that one can assign a color to each region of \mathcal{R} , using a total of at most k colors, such that the resulting coloring has the following property: For each point $p \in \bigcup_{R \in \mathcal{R}} R$ there is at least one region $R \in \mathcal{R}$ that contains p in its interior, whose color is *unique* among all regions in \mathcal{R} that contain p in their interior; in cellular network applications colors represent different communication frequencies and this is the scenario where a client (represented by a point p) is being “served” by a unique base station (represented by a region with a communication frequency not shared by any other base stations covering p).

Har-Peled and Smorodinsky [24] have shown that if \mathcal{R} is a set of n planar regions with union complexity $U(r)$ for every subset of size r , then $O(U(n) \log n/n)$ colors are enough to satisfy the CF-coloring property (here too, $U(\cdot)$ is assumed to be superadditive). Such a coloring can be computed in randomized expected time $O(U(n) \log n)$.

Applying the bounds in Theorems 4.10 and 4.11 once again, we obtain the following improved bounds on the CF-coloring problem for our fat objects:

Theorem 5.5. *Let \mathcal{R} be a set of n locally γ -fat objects in the plane. Then \mathcal{R} admits a conflict-free coloring with $O(\frac{1}{\gamma^4} 2^{c \log^* n} \log n)$ colors, where c is a constant depending only on the maximum complexity of the individual objects. Such a coloring can be computed in randomized expected time $O(\frac{n}{\gamma^4} 2^{c \log^* n} \log n)$. The bounds improve to $O((\log^* n + \frac{1}{\gamma} \log^2 \frac{1}{\gamma}) \log n)$ and $O(n(\log^* n + \frac{1}{\gamma} \log^2 \frac{1}{\gamma}) \log n)$, respectively, if the objects are γ -fat triangles.*

6 Concluding remarks

One of the new ingredients in our approach is that we reduce the problem to bounding the union complexity of a collection of towers. This is interesting for two reasons. First of all, towers are potentially easier to handle than (for example) triangles. Indeed, we gave a simple proof that, if the towers are all stabbed by a vertical line, then their union complexity is linear—such a result is not known for triangles. Second, the fatness of the towers in our reduction is lower bounded by an absolute constant, which gives a nice separation of concerns and allows us to obtain a bound for fat triangles where the dependency on γ only appears in the linear term (and not in the $O(n \log^* n)$ term).

Our bound on the complexity of the union of locally γ -fat objects is the first improvement over the previous bound of De Berg [7] and the first bound that breaks the “ $n \log n$ barrier.” Our bound on the complexity of the union of γ -fat triangles appears in two “installments”: it was $O(n2^{\alpha(n)} \log^* n)$ in the preliminary version of this paper [21], and constituted then the first improvement over the old bound in [30,31]. The present bound is an improvement of this bound as well as of an intermediate improvement of Pettie [33] in a follow-up study. There is still a (smaller) gap between these bounds and the known lower bound of $\Omega(n\alpha(n))$, which we conjecture to be the actual bound. Closing or even just tightening this gap remains a difficult and challenging problem.

Acknowledgments. We wish to thank Sariel Har-Peled and Jiří Matoušek for numerous helpful discussions.

References

- [1] P. K. Agarwal, J. Pach, and M. Sharir, State of the union (of geometric objects), In *Proc. Joint Summer Research Conf. on Discrete and Computational Geometry: 20 Years Later*, Contemp. Math. 452, Amer. Math. Society, pp. 9–48, 2008.
- [2] H. Alt, R. Fleischer, M. Kaufmann, K. Mehlhorn, S. Näher, S. Schirra and C. Uhrig, Approximate motion planning and the complexity of the boundary of the union of simple geometric figures, *Algorithmica* 8:391–406, 1992.
- [3] B. Aronov, E. Ezra, and M. Sharir, Small-size ε -nets for axis-parallel rectangles and boxes, *SIAM J. Comput.*, 39(7):3248–3282, 2010.
- [4] M. Bellare, S. Goldwasser, G. Lund, and A. Russell, Efficient probabilistically checkable proofs and applications to approximation, *Proc. 25th Annu. ACM Sympos. Theory Comput.*, pp. 294–304, 1993.
- [5] T. Chan, E. Grant, J. Koenemann, and M. Sharpe, Weighted capacitated, priority, and geometric set cover via improved quasi-uniform sampling, *Proc. 23rd Annu. ACM-SIAM Sympos. Discrete Algorithms*, pp. 1576–1585, 2012.
- [6] M. de Berg, Improved bounds on the union complexity of fat objects, *Discrete Comput. Geom.*, 40(1):127–140, 2008.
- [7] M. de Berg, Better bounds on the union complexity of locally fat objects, *Proc. 26th Annu. ACM Sympos. Comput. Geom.*, pp. 39–47, 2010.

- [8] M. de Berg and O. Schwarzkopf, Cutting and applications, *Int. J. Comput. Geom. Appl.*, 5:343–355, 1995.
- [9] M. de Berg, A. F. van der Stappen, J. Vleugels, M. J. Katz, Realistic input models for geometric algorithms, *Algorithmica*, 34(1):81–97, 2002.
- [10] H. Brönnimann and M. T. Goodrich, Almost optimal set covers in finite VC dimensions, *Discrete Comput. Geom.*, 14:463–479, 1995.
- [11] V. Chvátal, A greedy heuristic for the set-covering problem, *Math. Operations Res.*, 4(3):233–235, 1979.
- [12] K. L. Clarkson, Algorithms for polytope covering and approximation, *Proc. 3rd Workshop Algor. Data Struct.*, pp. 246–252, 1993.
- [13] K. L. Clarkson and P. W. Shor, Applications of random sampling in computational geometry, II, *Discrete Comput. Geom.*, 4:387–421, 1989.
- [14] K. L. Clarkson and K. Varadarajan, Improved approximation algorithms for geometric set cover, *Discrete Comput. Geom.*, 37:43–58, 2007.
- [15] H. Edelsbrunner, L. J. Guibas, and M. Sharir, The complexity and construction of many faces in arrangements of lines and of segments, *Discrete Comput. Geom.*, 5(2):161–196, 1990.
- [16] A. Efrat, The complexity of the union of (α, β) -covered objects, *SIAM J. Comput.*, 34:775–787, 2005.
- [17] A. Efrat and M. Katz, On the union of κ -curved objects, *Comput. Geom. Theory Appl.*, 14:241–254, 1999.
- [18] A. Efrat, G. Rote, and M. Sharir, On the union of fat wedges and separating a collection of segments by a line, *Comput. Geom. Theory Appl.*, 3:277–288, 1993.
- [19] A. Efrat and M. Sharir, The complexity of the union of fat objects in the plane, *Discrete Comput. Geom.*, 23:171–189, 2000.
- [20] G. Even, D. Rawitz, and S. Shahar, Hitting sets when the VC-dimension is small, *Inform. Process. Lett.*, 95:358–362, 2005.
- [21] E. Ezra, B. Aronov, and M. Sharir, Improved bound for the union of fat triangles, *Proc. 22nd Annu. ACM-SIAM Sympos. Discrete Algorithms*, pp. 1778–1785, 2011.
- [22] E. Ezra and M. Sharir, Almost tight bound for the union of fat tetrahedra in three dimensions, *J. ACM* 57(1), Article No. 2, 2009.
- [23] U. Feige, A threshold of $\ln n$ for approximating set cover, *J. ACM*, 45(4):634–652, 1998.
- [24] S. Har-Peled and S. Smorodinsky, Conflict-free coloring of points and simple regions in the plane, *Discrete Comput. Geom.*, 34(1):47–70, 2005.
- [25] R. Hartshorne. *Algebraic Geometry*, Springer-Verlag, New York 1977.

- [26] D. Haussler and E. Welzl, ε -nets and simplex range queries, *Discrete Comput. Geom.*, 2:127–151, 1987.
- [27] M. J. Katz, M. H. Overmars, and M. Sharir, Efficient hidden surface removal for objects with small union size, *Comput. Geom. Theory Appl.*, 2(4):223–234, 1992.
- [28] K. Kedem, R. Livne, J. Pach, and M. Sharir, On the union of Jordan regions and collision-free translational motion amidst polygonal obstacles, *Discrete Comput. Geom.*, 1:59–71, 1986.
- [29] M. van Kreveld, On fat partitioning, fat covering, and the union size of polygons, *Comput. Geom. Theory Appl.*, 9(4):197–210, 1998.
- [30] J. Matoušek, N. Miller, J. Pach, M. Sharir, S. Sifrony, and E. Welzl, Fat triangles determine linearly many holes, *Proc. 32nd Annu. IEEE Sympos. Found. Comput. Sci.*, pp. 49–58, 1991.
- [31] J. Matoušek, J. Pach, M. Sharir, S. Sifrony, and E. Welzl, Fat triangles determine linearly many holes, *SIAM J. Comput.*, 23:154–169, 1994.
- [32] J. Pach and G. Tardos, On the boundary complexity of the union of fat triangles, *SIAM J. Comput.*, 31:1745–1760, 2002.
- [33] S. Pettie, On the structure and composition of forbidden sequences, with geometric applications, *Proc. 27th Annu. ACM Sympos. Comput. Geom.*, pp. 370–379, 2011.
- [34] M. Sharir and P. K. Agarwal, *Davenport-Schinzel Sequences and Their Geometric Applications*, Cambridge University Press, New York, 1995.
- [35] A. F. van der Stappen, *Motion Planning amidst Fat Obstacles*, PhD Dissertation, Dept. Comput. Sci., Utrecht Univ., Utrecht, Netherlands, 1994.
- [36] K. Varadarajan, Epsilon nets and union complexity, *Proc. 25th Annu. ACM Sympos. Comput. Geom.*, pp. 11–16, 2009.
- [37] K. Varadarajan, Weighted geometric set cover via quasi-uniform sampling, *Proc. 42th Annu. ACM Sympos. Theory Comput.*, pp. 641–648, 2010.