Discrete Comput Geom (2011) 45: 45–64
DOI 10.1007/s00454-010-9312-x

On the Union of Cylinders in Three Dimensions

Esther Ezra

Received: 21 January 2009 / Revised: 23 October 2010 / Accepted: 16 November 2010 /
Published online: 21 December 2010

Abstract We show that the combinatorial complexity of the union of n infinite cylinders in $\mathbb{R}^3$, having arbitrary radii, is $O(n^{2+\varepsilon})$, for any $\varepsilon > 0$; the bound is almost tight in the worst case, thus settling a conjecture of Agarwal and Sharir (Discrete Comput. Geom. 24:645–685, 2000), who established a nearly-quadratic bound for the restricted case of *nearly congruent cylinders*. Our result extends, in a significant way, the result of Agarwal and Sharir (Discrete Comput. Geom. 24:645–685, 2000), in particular, a simple specialization of our analysis to the case of nearly congruent cylinders yields a nearly-quadratic bound on the complexity of the union in that case, thus significantly simplifying the analysis in Agarwal and Sharir (Discrete Comput. Geom. 24:645–685, 2000). Finally, we extend our technique to the case of "cigars" of arbitrary radii (that is, Minkowski sums of line-segments and balls) and show that the combinatorial complexity of the union in this case is nearly-quadratic as well. This problem has been studied in Agarwal and Sharir (Discrete Comput. Geom. 24:645–685, 2000) for the restricted case where all cigars have (nearly) equal radii. Based on our new approach, the proof follows almost verbatim from the analysis for infinite cylinders and is significantly simpler than the proof presented in Agarwal and Sharir (Discrete Comput. Geom. 24:645–685, 2000).

Keywords Union of simply-shaped bodies · Geometric arrangements · Lower envelopes · $(1/r)$-cuttings

Part of the work on this paper was performed at the Department of Computer Science, Duke University, Durham, NC 27708-0129, USA, and has been supported by NSF under grants CNS-05-40347, CFF-06-35000, and DEB-04-25465, by ARO grants W911NF-04-1-0278 and W911NF-07-1-0376, and by an NIH grant 1P50-GM-08183-01. A preliminary version of this paper has appeared in *Proc. 49th Annu. IEEE Sympos. Found. Comput. Sci.* 2008.

E. Ezra (✉)
Courant Institute of Mathematical Sciences, New York University, New York, NY 10012, USA
e-mail: esther@courant.nyu.edu

1 Introduction

Let $\mathcal{K} = \{K_1, \ldots, K_n\}$ be a collection of n infinite cylinders in $\mathbb{R}^3$, and put $\mathcal{C} = \{\partial K_1, \ldots, \partial K_n\}$, where ∂O is the boundary of an object O; with a slight abuse of notation, we also refer to the elements of $\mathcal{C}$ as cylinders, but sometimes use the term "cylinder boundary" to refer to a boundary of a cylinder when it is not sufficiently clear from the context. Let $\mathcal{A}(\mathcal{C})$ denote the three-dimensional arrangement induced by the cylinders in $\mathcal{C}$, i.e., the decomposition of 3-space into vertices, edges, faces, and three-dimensional cells, each being a maximal connected set contained in the intersection of a fixed subcollection (possibly empty) of the cylinders of $\mathcal{C}$ and not meeting any other cylinder. Let $\mathcal{U} = \bigcup_{i=1}^{n} K_i$ denote the union of $\mathcal{K}$. The *combinatorial complexity* of $\mathcal{U}$ is the number of vertices, edges, and faces of the arrangement $\mathcal{A}(\mathcal{C})$ appearing on the boundary of the union. The problem studied in this paper is to obtain a nearly-quadratic upper bound (in the number of cylinders) on the combinatorial complexity of $\mathcal{U}$.

The case of cylinders of equal radii arises in the motion planning problem of a ball amid line obstacles in $\mathbb{R}^3$. If the obstacles are n pairwise disjoint cylinders, then the complement of the union of their Minkowski sums with the moving ball is the free space for the center of the ball. Since these sums are possibly intersecting cylinders of nonequal radii (where each cylinder radius is the sum of the corresponding obstacle radius and the ball radius), our result implies that the combinatorial complexity of the free space is $O(n^{2+\varepsilon})$, for any $\varepsilon > 0$. Note that the problem studied in this paper is more general, in the sense that it extends to any set of (possibly intersecting) cylinders, and not only to settings obtained by the motion planning problem.

Previous results The problem of determining the combinatorial complexity of the union of geometric objects has received considerable attention in the past twenty years, although most of the earlier work has concentrated on the planar case. See [1] for a recent comprehensive survey of the area.

The case involving *pseudodiscs* (that is, a collection of simply connected planar regions, where the boundaries of any two distinct objects intersect at most twice) arises for Minkowski sums of a fixed convex object with a set of pairwise disjoint convex objects (which is the problem one faces in translational motion planning of a convex robot) and has been studied by Kedem et al. [21]. In this case, the union has only linear complexity in the total input size. Matoušek et al. [23, 24] proved that the union of n δ-*fat* triangles (where a triangle is δ-fat if each of its angles is at least δ) in the plane has only $O(n)$ holes, and its combinatorial complexity is $O(n \log\log n)$. The constant of proportionality, which depends on the fatness factor δ, has later been improved by Pach and Tardos [27]. Very recently Ezra et al. [15] have improved the $O(n \log\log n)$ bound to $O(n 2^{\alpha(n)} \log^* n)$, where $\alpha(\cdot)$ is the extremely slowly growing inverse Ackermann function [28]. Extending the study to the realm of curved objects, Efrat and Sharir [14] studied the union of planar convex fat objects. Here we say that a planar convex object c is γ-*fat*, for some fixed $\gamma > 1$, if there exist two concentric disks, $D \subseteq c \subseteq D'$, such that the ratio between the radii of D' and D is at most γ. In this case, the combinatorial complexity of the union of n such objects, such that the boundaries of each pair of objects intersect in a constant number of points, is $O(n^{1+\varepsilon})$, for any $\varepsilon > 0$. See also Efrat and Katz [13], Efrat [12], and de Berg [10, 11] for related (and slightly sharper) nearly-linear bounds in n.

In three and higher dimensions, it was shown by Aronov et al. [6] that the complexity of the union of k convex polyhedra with a total of n facets in $\mathbb{R}^3$ is $O(k^3 + nk \log k)$, and it can be $\Omega(k^3 + nk\alpha(k))$ in the worst case. The bound was improved by Aronov and Sharir [5] to $O(nk \log k)$ (and $\Omega(nk\alpha(k))$) when the given polyhedra are Minkowski sums of a fixed convex polyhedron with k pairwise-disjoint convex polyhedra. (This problem arises in the case of a translating convex polyhedral robot in $\mathbb{R}^3$ amid a collection of polyhedral obstacles.) Boissonnat et al. [7] proved that the maximum complexity of the union of n axis-parallel hypercubes in $\mathbb{R}^d$ is $\Theta(n^{\lceil d/2 \rceil})$, and that the bound improves to $\Theta(n^{\lfloor d/2 \rfloor})$ if all hypercubes have the same size. Pach et al. [26] showed that the combinatorial complexity of the union of n nearly congruent arbitrarily oriented cubes in three dimensions is $O(n^{2+\varepsilon})$, for any $\varepsilon > 0$ (see also [25] for a subcubic bound on the complexity of the union of fat wedges in 3-space). Agarwal and Sharir [3] have shown that the complexity of the union of n congruent infinite cylinders is $O(n^{2+\varepsilon})$, for any $\varepsilon > 0$. In fact, the more general problem studied in [3] involves the union of the Minkowski sums of n pairwise disjoint triangles with a ball (where congruent infinite cylinders are obtained when the triangles become lines), and the $O(n^{2+\varepsilon})$ bound is extended in [3] to this case as well. Aronov et al. [4] showed that the union complexity of n *κ-round* objects in $\mathbb{R}^3$ is $O(n^{2+\varepsilon})$, for any $\varepsilon > 0$, where an object c is κ-round if for each $p \in \partial c$, there exists a ball $B \subset c$ that touches p and its radius is at least $\kappa \cdot \mathrm{diam}(c)$. The bound is $O(n^{3+\varepsilon})$, for any $\varepsilon > 0$, for κ-round objects in $\mathbb{R}^4$. Finally, Ezra and Sharir [16] have recently shown that the complexity of the union of n *δ-fat* tetrahedra (that is, tetrahedra, each of whose four solid angles at its four respective apices is at least δ) of arbitrary sizes in $\mathbb{R}^3$ is $O(n^{2+\varepsilon})$, for any $\varepsilon > 0$. This result immediately yields a nearly-quadratic bound on the complexity of the union of arbitrary cubes and thus generalizes the result of Pach et al. [26], who showed this bound only for the case where the cubes have nearly equal edge lengths. Each of the above known nearly-quadratic bounds (for the three-dimensional case) is almost tight in the worst case.

The above results indicate that the combinatorial complexity of the union of the input bodies in these cases is roughly "one order of magnitude" smaller than the complexity of the full arrangement of their boundaries. While considerable progress has been made on the analysis of unions in three dimensions, the case of the union of infinite cylinders of arbitrary radii has so far been remained elusive.

We also note that when the complement of the union of the cylinders is *connected*, then one can apply the $O(n^{2+\varepsilon})$ bound established by Halperin and Sharir [19] (see also [18]) on the boundary of a single cell in an arrangement of cylinders (in fact, these results are more general and can also be applied to other objects satisfying certain properties). However, when the complement of the union is not connected, these techniques do not guarantee the aforementioned bound, and we thus need to resort to a different approach.

Our result In this paper we make a significant progress on the problem of bounding the complexity of the union of infinite cylinders of arbitrary radii in 3-space and show a nearly-quadratic bound (in the input size) on this complexity, thus settling a conjecture of Agarwal and Sharir [3], who showed this bound only for the case where the cylinders are (nearly) congruent. Our bound, which is the first known nontrivial bound for this general problem, is almost tight in the worst case.

Specifically, we show that the complexity of the union of n cylinders as above is $O(n^{2+\varepsilon})$, for any $\varepsilon > 0$, where the constant of proportionality depends on ε. The analysis is based on some of the ideas presented in [3, 16], and on a divide-and-conquer mechanism, which partitions space into triangular-prism subcells, so that, on average, the overwhelming majority of the cylinders intersecting a subcell Δ are "good", in the sense that they behave as *functions* within Δ with respect to some direction ρ. Then, using a fundamental property shown in [2, 22], we conclude that the complexity of the union of the cylinders, bounded by these functions (in the positive or negative ρ-direction), is only nearly-quadratic in the total number of functions. It then only remains to analyze the number of other types of vertices (incident to some of the few "bad" cylinders that intersect Δ), a task which is handled by the divide-and-conquer mechanism (see below for details).

The problem studied in this paper is a generalization of the case where all the cylinders are of equal radii—a problem that has been studied by Agarwal and Sharir [3]. We show that a simple specialization of our analysis to that case yields the same asymptotic bound on the complexity of the union as above. Note that we use a variant of some of the ideas given in [3]; however, most of the analysis steps taken in that previous study are no longer needed.

We extend our analysis to the case of "cigars" of arbitrary radii, that is, Minkowski sums of line-segments and balls, and show that the bound on the combinatorial complexity of the union is nearly-quadratic in n in this case as well. This problem has been studied in the previous work by Agarwal and Sharir [3] for the restricted case where the balls in the Minkowski sum are congruent. Here too, our analysis is significantly simpler than that of the previous study, and, in particular, the original problem is much easier to extend to this case using our new approach.

We note that our approach, as well as the preceding approach presented in [3], is based on the decomposition of space into vertical prisms and then observing that the overall majority of the union vertices in each prism cell appear on some monotone surface (with respect to some direction ρ). Nevertheless, the decomposition in the previous study uses a uniform grid, which, with an additional intricate machinery, can only handle congruent cylinders, whereas our approach exploits a more sophisticated divide-and-conquer mechanism in order to classify the cylinders as either good (that is, behave as functions within a cell of the decomposition) or bad (otherwise). This enables us to capture cylinders of arbitrary radii, and thus we need not rely on the additional machinery of [3]. We note that our divide-and-conquer mechanism is somewhat inspired by the approach of [16]. Nevertheless, the decomposition presented there is different and more intricate than the one used here, which, by combining our new observations, yields the aforementioned nearly-optimal bounds—see below.

2 The Complexity of the Union

2.1 Preliminaries and Overview

We first assume that the cylinders in $\mathcal{K}$ are in general position, that is, no two of them have parallel or intersecting axes, no two cylinders are tangent to each other, no curve of intersection of the boundaries of any two cylinders is tangent to a third one, and no four cylinder boundaries meet. Following the arguments in [3], this assumption involves no loss of generality. This general position assumption implies that each vertex of the arrangement of the cylinders lies on exactly three cylinder boundaries and is thus incident upon only a constant number of edges and faces. The number of edges and faces on $\partial\mathcal{U}$ that are not incident upon any vertex is $O(n^2)$. Thus the combinatorial complexity of $\mathcal{U}$ is $O(n^2 + |V(\mathcal{C})|)$, where $V(\mathcal{C})$ is the set of vertices of $\mathcal{A}(\mathcal{C})$ that appear on ∂U. Our main result is:

Theorem 2.1 *For any set $\mathcal{C}$ of n infinite cylinders in $\mathbb{R}^3$, $|V(\mathcal{C})| = O(n^{2+\varepsilon})$, where the constant of proportionality depends on ε. The bound is almost tight in the worst case.*

It is relatively easy (using standard techniques; see, e.g., [28]) to construct a set of n infinite cylinders that yield $\Omega(n^2)$ vertices on the boundary of their union (see Fig. 4(a) and [3] for further details). We thus devote the remainder of this section to deriving the upper bound stated in Theorem 2.1.

We note that it is crucial to assume that the cylinders are infinite. Otherwise, the combinatorial complexity of their union is $\Omega(n^3)$ in the worst case. Indeed, suppose we have a set of n cylinders, each of which with a sufficiently large radius and height that is arbitrarily close to 0. We can now arrange these cylinders in a (three-dimensional) grid-like structure, resulting in $\Omega(n^3)$ holes in the union; see Fig. 1(a).

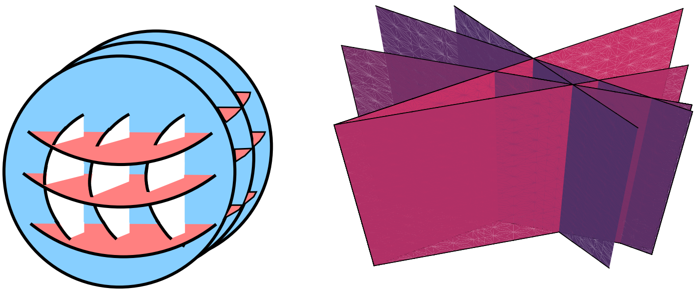

(a) (b)

Fig. 1 (**a**) Coins in 3-space arranged in a grid-like form. (**b**) The decomposition into vertical prism cells

2.2 The Problem Decomposition—An Overview

We use a divide-and-conquer approach, based on $(1/r)$-cuttings [8, 9, 20]. Given a collection $\mathcal{R}$ of n surface patches in $\mathbb{R}^d$ of constant description complexity,[1] a $(1/r)$-cutting for $\mathcal{R}$ is the subdivision of space into *pseudo-simplicial subcells* (that is, each subcell is not necessarily piecewise linear but has the topological structure of a simplex) such that every cell is intersected by at most n/r elements of $\mathcal{R}$, where $1 \le r \le n$ is the parameter of the cutting. We call this property the *cutting property*. By the random sampling theory of Clarkson and Shor [8, 9], as well as the epsilon-net theorem of Haussler and Welzl [20], a $(1/r)$-cutting can be constructed by taking a random sample of $O(r \log r)$ (with a sufficiently large constant of proportionality) of the elements in $\mathcal{R}$, constructing their arrangement, and then applying a simplicial decomposition of its cells. The cutting property is then guaranteed to be satisfied with constant probability. In what follows, we rely on the existence of $(1/r)$-cuttings of the above size and just fix a sample of $O(r \log r)$ elements that satisfy the cutting property. In some cases we use a more intricate decomposition than just that of considering the entire arrangement; nevertheless, the cutting property will still hold in these cases—see below.

Specifically, in our problem we project all the cylinders in $\mathcal{K}$ onto the xy-plane, thereby obtaining a set $\mathcal{K}^0$ of n infinite strips. Let $\mathcal{L}$ denote the set of their bounding lines. (We assume that the coordinate system is generic, so none of these lines projects to a single point.)

We construct a $(1/r_0)$-cutting of the planar arrangement $\mathcal{A}(L)$, by taking a sample R_0 (satisfying the cutting property) of $O(r_0 \log r_0)$ lines of $\mathcal{L}$, for some sufficiently large constant parameter r_0 to be fixed shortly, constructing the planar arrangement $\mathcal{A}(R_0)$, and triangulating each of its cells, using, e.g., bottom-vertex triangulation. The number of simplices is proportional to the overall complexity of $\mathcal{A}(R_0)$ and is thus $O({r_0}^2 \log^2 r_0)$. By the cutting property, each simplex of the resulting decomposition is intersected by at most n/r_0 lines of $\mathcal{L}$.

We next lift the subtriangles of each of the cells of $\mathcal{A}(R_0)$ into vertical prisms. Let Ξ denote the collection of these prism-cells; see Fig. 1(b) for an example.

Our goal is to bound the number of intersection vertices of the union in each cell of Ξ separately and then sum these bounds over all cells. Fix a cell Δ of Ξ. We classify each cylinder $K \in \mathcal{K}$ that intersects Δ as being either *wide* in Δ if the radius r of K satisfies $r \ge w/2$, where w is the *width* of Δ, or *narrow* otherwise.[2] See Fig. 2 for an example.

As a consequence, each intersection vertex v of the union that appears in Δ is classified as either *WWW* if all three cylinders that are incident to v are wide in Δ, *WWN* if two of them are wide and one is narrow, *WNN* if one of them is wide and two are narrow, or *NNN* if all of them are narrow in Δ. In all four cases, the three relevant cylinders are distinct.

[1] A set has *constant description complexity* if it is a semi-algebraic set defined by a constant number of polynomial equations and inequalities of constant maximum degree.

[2] Here we slightly deviate from the standard width definition and define it to be the minimum distance between any pair of parallel vertical planes that bound Δ.

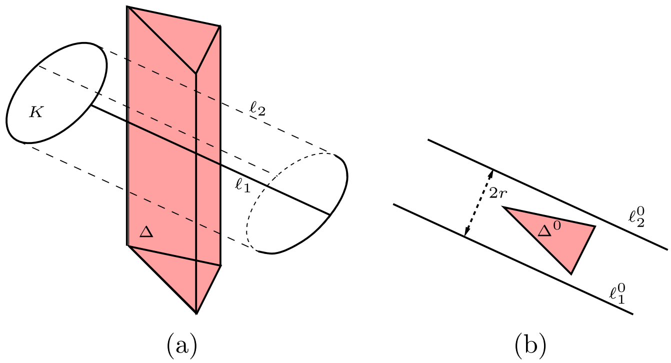

Fig. 2 The cylinder K is wide within the prism-cell Δ. (**a**) The pair of the silhouette lines ℓ_1, ℓ_2 of K do not meet Δ. (**b**) The projection of this scene onto the xy-plane. The lines ℓ_1^0, ℓ_2^0 are the respective projections of ℓ_1, ℓ_2, Δ^0 is the projection of Δ, and r is the radius of K. The strip bounded by ℓ_1^0, ℓ_2^0 contains Δ^0 in its interior

Let Δ be a prism-cell of Ξ. We first observe that if the radius r of a cylinder K that meets Δ is smaller than $w/2$ (that is, K is narrow in Δ), at least one of the two *silhouette* lines of K (this is the locus of all the vertical tangency points on the boundary of K) must intersect Δ. Indeed, if both lines do not meet Δ, the diameter $d := 2r$ of K must be larger than w, as is easily verified, contrary to the assumption that K is narrow in Δ. We thus charge the crossing of K and Δ to that of ℓ and Δ, and since $\mathcal{L}$ is in fact the set of the projections of the silhouette lines onto the xy-plane, we conclude that Δ meets at most $\frac{n}{r_0}$ narrow cylinders; see once again Fig. 2.

In what follows, for every prism-cell $\Delta \in \Xi$, let $\mathcal{W} = \mathcal{W}^\Delta$ (resp., $\mathcal{N} = \mathcal{N}^\Delta$) be the set of the wide (resp., narrow) cylinders within Δ. Put $M_W = M_W^\Delta := |\mathcal{W}^\Delta|$ and $M_N = M_N^\Delta := |\mathcal{N}^\Delta|$. As just discussed, $M_W \le n$, $M_N \le \frac{n}{r_0}$.

2.3 The Overall Recursive Analysis

We now present a recursive decomposition, where in each step, we are given a (parent) cell Δ_0 and a set $\mathcal{N}^{\Delta_0}$ of narrow cylinders that meet Δ_0, and we construct $O(r_0{}^2 \log^2 r_0)$ prism-subcells Δ of Δ_0 as above. At the first recursive step, Δ_0 equals to the entire three-dimensional space, and $\mathcal{N}^{\Delta_0}$ is the entire set of the input cylinders (and thus $|\mathcal{N}^{\Delta_0}| = n$ at that step). As the space is progressively cut up into subcells (which follows by the planar partition that we apply), more and more narrow cylinders become wide, and the size of the set $\mathcal{N}^{\Delta_0}$ keeps decreasing. During each step of the recursion, we immediately dispose of any *new* WWW-, WWN-, and WNN-vertices within each subcell Δ (that is, vertices that were NNN in the parent cell Δ_0) and continue with bounding the number of the remaining NNN-vertices recursively. In particular, this implies that we dispose of all the *new* wide cylinders (that is, cylinders that were narrow in Δ_0) in a single step and process only the (remaining) narrow ones. The recursion bottoms out when $M_N \le c$, for some absolute constant $c \ge 3$. In this case the number of the remaining intersection vertices of the union within the current cell Δ_0 is $O(c^3) = O(1)$. We show below:

Lemma 2.2 *The overall number of WWW-, WWN-, and WNN-vertices in a subcell* Δ *is* $O((M_N^\Delta + M_W^\Delta)^{2+\delta})$, *for any* $\delta > 0$.

The bound in Lemma 2.2 implies that the overall number of new WWW-, WWN-, and WNN-vertices generated within each subcell Δ of Δ_0 is $O((M_N^{\Delta_0})^{2+\delta})$ (since $M_W^\Delta \le M_N^{\Delta_0}$, $M_N^\Delta \le \frac{M_N^{\Delta_0}}{r_0}$), for a total of $O(r_0^2 \log^2 r_0 (M_N^{\Delta_0})^{2+\delta}) = O((M_N^{\Delta_0})^{2+\delta})$ such vertices (recall that r_0 is constant), over all subcells of Δ_0. Let $U_0(M_N)$ denote the maximum number of intersection vertices that appear on the boundary of the union at a recursive step involving up to M_N narrow cylinders. Then U_0 satisfies the following recurrence:

$$U_0(M_N) \le \begin{cases} O(M_N^{2+\delta}) + O(r_0^2 \log^2 r_0) U_0(M_N/r_0) & \text{if } M_N > c, \\ O(1) & \text{if } M_N \le c, \end{cases}$$

where $\delta > 0$ is arbitrary, $c \ge 3$ is an appropriate constant as above, and the constant of proportionality in the nonrecursive term depends on r_0 (and on δ).

To solve the recurrence, for a given $\varepsilon > 0$, we substitute $r_0 = M_N^{\varepsilon'}$, for $\varepsilon' = o(\varepsilon)$. Then, using induction on M_N, it is easy to verify that the solution of this recurrence is $U_0(M_N) = O(M_N^{2+\varepsilon})$, for any $\varepsilon > 0$, slightly larger than δ in the nonrecursive term, but still arbitrarily close to 0, with a constant of proportionality that depends on ε (see [17] for similar considerations). Substituting the initial value $M_N = n$, we conclude that the overall number of vertices of the union is $O(n^{2+\varepsilon})$, for any $\varepsilon > 0$, as asserted.

We thus devote the remainder of this section to deriving the bound, stated in Lemma 2.2, on the number of vertices of the first three types, which is the major part of the analysis.

The number of WWW-vertices of the union Our next goal is to show that the number of WWW-vertices of the union, contained in the fixed prism-cell Δ, is only nearly-quadratic in M_W^Δ.

Let H, H' be the pair of (parallel and vertical) planes that determine the width w of Δ. We use a simple variant of the analysis given in [3], in order to show that the WWW-vertices of Δ (in fact, it is sufficient to assume that they lie in the slab bounded by H, H') appear on the boundary of the region enclosed between a lower envelope of a collection of portions of the wide cylinders and an upper envelope of another such collection (this is the so-called *sandwich region*). The analysis of [2] (see also [22]) implies that this complexity is only nearly-quadratic in the total number of cylinder portions.

In the sequel, we follow almost verbatim the analysis in [3] and repeat some of the details there for the sake of completeness. The main difference lies in the cutting-based decomposition that we use, and, in particular, in the pair of planes H, H' for each cell Δ in the decomposition, which have a crucial role in the analysis, whereas the decomposition in [3] consists of just a grid of infinite square-prisms whose axes are parallel to some fixed direction u. The main ideas that we borrow from [3] are (i) the partition of the boundary of a wide cylinder into "canonical strips" τ and

(ii) the definition of a "good direction" ρ (modified to be matched with our decomposition) with respect to a strip and to a vertex v of the union (which appears on the relative interior of a triple of canonical strips). Then, based on these definitions, we observe that once we *enter* into a cylinder K, bounded by a strip τ, through a WWW-vertex $v \in \tau \cap \Delta$ of the union in the ρ-direction (that is, we draw a ray in the ρ-direction whose endpoint is at v), we make a sufficiently long way inside K with respect to the width of Δ (in fact, a more precise description is given in Lemma 2.4). Then the property that v appears on the sandwich region enclosed between an upper and a lower envelope of canonical strips (appropriately defined at the ρ-direction) follows almost immediately—see below.

Let κ be a sufficiently large constant, whose value will be fixed shortly. We partition each of the wide cylinder boundaries into κ *canonical strips* (that is, portions of the cylinder boundary, each of which is bounded by a pair of generator lines, parallel to the cylinder axis), each with an angular span of $2\pi/\kappa$ in the cylindrical coordinate frame induced by the cylinder.

Let $\mathbb{S}^2$ be the unit sphere of directions in $\mathbb{R}^3$. We say that a direction $\rho \in \mathbb{S}^2$ is *good* for a strip τ if the following two conditions hold.

(i) The angle between ρ and the (outer) normal $\mathbf{n}$ of either of the planes H, H' is smaller than $\pi/2 - \pi/\kappa$.
(ii) Each line tangent to (the relative interior of) τ forms an angle of at least π/κ with ρ.

In other words, the first condition implies that the inner product $\rho \cdot \mathbf{n}$ is sufficiently large (in particular, they are not orthogonal), and the second condition implies that when we enter into the cylinder K, which contains τ as a bounding strip, in the ρ-direction from a point on τ, we do not remain close to the boundary of K (and, in particular, to τ). Intuitively, these two conditions imply that when we move in the ρ-direction, we expect to leave the strip enclosed between H, H' relatively soon, but make a sufficiently long way inside K.

We say that ρ is a *good direction* for a vertex v of the union, incident upon three canonical strips, if it is good for each of these strips; see Fig. 3. Using similar considerations as in [3], the set B_τ of bad directions for a fixed strip τ is the union $B_1 \cup B_2$, where

(a) B_1 is contained in a spherical band consisting of all points lying at spherical distance at most π/κ from the great circle on $\mathbb{S}^2$ normal to $\mathbf{n}$ (in fact, this circle is obtained by intersecting $\mathbb{S}^2$ with a plane parallel to H (or H') through its center).
(b) B_2 is contained in a spherical band consisting of all points lying at spherical distance at most $2\pi/\kappa$ from a great circle on $\mathbb{S}^2$; see [3] for the easy proof.

Claim 2.3 (Agarwal and Sharir [3]) *The following properties hold*:

1. *The area of* B_1 *is* $4\pi \sin(\pi/\kappa)$.
2. *The area of* B_2 *is* $4\pi \sin(2\pi/\kappa)$.
3. *The area of the set of good directions for* v *is at least* $4\pi[1 - 3\sin(2\pi/\kappa) - \sin(\pi/\kappa)]$.

Fig. 3 The vector ρ is a good direction for the vertex v of the union, incident upon three canonical strips τ_1, τ_2, τ_3, where $\mathbf{n}$ is the normal to the planes H, H'

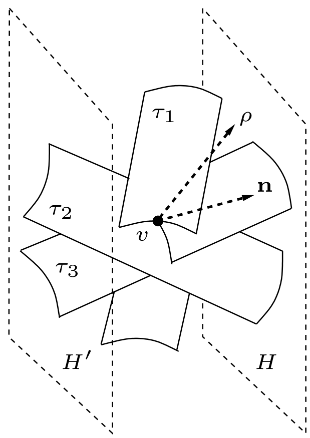

As observed in [3], the set of the good directions contains a spherical cap of some constant opening angle δ if κ is sufficiently large. It then implies that there exists a set $\mathcal{Z}$ of $O(1/\delta^2)$ points on $\mathbb{S}^2$ such that each vertex v on the boundary of the union has at least one good direction taken from $\mathcal{Z}$, for some sufficiently small value of δ that depends on κ. Note that the above considerations do not assume that v is a vertex of type WWW. For each fixed $\rho \in \mathcal{Z}$, let $V_\Delta(\rho)$ be the subset of all vertices (of any of the above four types) of the union that lie inside Δ, for which ρ is a good direction.

Let σ be the slab enclosed between H, H'. We now partition σ (and thus Δ) into congruent subslabs σ' by adding $t > \frac{1}{\sin^2(\pi/\kappa)}$ planes parallel to H inside σ at equal distances. Our next goal is to show that the strips τ for which ρ is a good direction behave as functions (with respect to that direction) within σ'. The following lemma is a simple variant of [3, Lemma 2.9]:

Lemma 2.4 (Agarwal and Sharir [3]) *Let σ' be a subslab of σ, and let v be a WWW-vertex of the union, contained in $\sigma' \cap \Delta$ and incident upon three strips τ_1, τ_2, τ_3. Let $\rho \in \mathcal{Z}$ be a good direction for v. Then any line parallel to ρ intersects τ_1 in at most one point. Moreover, for any point $u \in \tau_1 \cap \sigma'$, in which we enter into the cylinder K_1 bounded by τ_1 in the direction parallel to ρ, we reach $\partial\sigma'$ before exiting K_1. Similar properties hold for τ_2, τ_3.*

Proof Suppose first that there is a line parallel to ρ that intersects τ_1 twice. This would imply that τ_1 must have a tangent in the ρ-direction, which can be obtained by just translating this line until it exits K_1 (through a point on τ_1), contrary to the assumption that ρ is a good direction.

As to the second assertion of the lemma, let u' be the other intersection point of ∂K_1 and the line passing through u and parallel to ρ. It is easy to verify that $|uu'|$ is minimized when uu' is orthogonal to the axis of K_1 and forms an angle π/κ with the tangent plane to K_1 at u (this latter property is determined by condition (ii) that ρ satisfies as a good direction). It thus follows that $|uu'| \geq 2r_1 \sin(\pi/\kappa)$, where r_1 is the radius of K_1. On the other hand, since uu' forms an angle of at most $\pi/2 - \pi/\kappa$ with the normal $\mathbf{n}$ (to either of the two bounding planes of σ'), it follows that the length of the projection of uu' on $\mathbf{n}$ is at least $|uu'|\sin(\pi/\kappa) \geq 2r_1 \sin^2(\pi/\kappa) \geq w \sin^2(\pi/\kappa)$, where w is the width of σ (that is, the distance between H, H'). Thus, if we choose $t > \frac{1}{\sin^2(\pi/\kappa)}$, u' must lie outside σ'. This completes the proof of the lemma. □

We now denote by $\mathcal{T}_{\sigma'}(\rho)$ the set of canonical strips τ that intersect the subslab σ' and have at least one vertex in $V_\Delta(\rho)$. We clip all the strips to $\Delta \cap \sigma'$. A (clipped) strip τ in $\mathcal{T}_{\sigma'}(\rho)$ is ρ-*upper* (resp., ρ-*lower*) if, for any point $u \in \tau$, the point $u + \alpha\rho$ lies in the exterior (resp., interior) of the cylinder K whose boundary contains τ, for sufficiently small values of $\alpha > 0$. Let $\mathcal{T}^+_{\sigma'}(\rho)$ (resp., $\mathcal{T}^-_{\sigma'}(\rho)$) be the set of the ρ-upper (resp., ρ-lower) strips of $\mathcal{T}_{\sigma'}(\rho)$. The ρ-*upper envelope* of $\mathcal{T}^+_{\sigma'}(\rho)$ is the set of points u on the strips of $\mathcal{T}^+_{\sigma'}(\rho)$ such that a ray from u in the $(+\rho)$-direction does not meet any other strip in $\mathcal{T}^+_{\sigma'}(\rho)$. The ρ-*lower envelope* of $\mathcal{T}^-_{\sigma'}(\rho)$ is defined analogously.

Following the analysis of [3] and Lemma 2.4, each WWW-vertex $v \in V_\Delta(\rho)$ must lie on the boundary of the sandwich region enclosed between the ρ-upper envelope of the strips in $\mathcal{T}^+_{\sigma'}(\rho)$ and the ρ-lower envelope of the strips in $\mathcal{T}^-_{\sigma'}(\rho)$, and, according to the results of [2, 22], the number of these vertices v is $O(M_W^{2+\varepsilon})$, for any $\varepsilon > 0$, with a constant of proportionality that depends on κ and ε. Repeating the analysis for all subslabs σ' and all good directions $\rho \in \mathcal{Z}$, we obtain a similar bound on the overall number of WWW-vertices of the union that are contained in Δ.

Remark We note that it is crucial to process only the strips that have at least one vertex v of the union for which ρ is a good direction and construct the ρ-lower and ρ-upper envelopes only with respect to these strips. All the remaining strips should not be considered at that step, since ρ is not necessarily a good direction for them, and they thus may not behave as graphs of bivariate functions at that direction (within σ'). In particular, we ignore wide cylinders whose corresponding strips do not contain any such vertices v.

The number of WWN-vertices of the union We next establish a nearly-quadratic bound (to be specified below) on the number of WWN-vertices of the union, with respect to a fixed prism-cell Δ. Let us first fix a good direction $\rho \in \mathcal{Z}$ and a subslab σ'. We bound the number of WWN-vertices in each subcell $\Delta' = \Delta \cap \sigma'$ separately and then sum these bounds over all these subcells. Note that some of the narrow cylinders within Δ may become wide in Δ', nevertheless, we regard them as narrow; this does not effect the analysis. We thus have, at each step of the analysis, a prism-cell Δ' (confined to a slab σ'), a set $\mathcal{W} = \mathcal{W}^{\Delta'}$ of M_W wide cylinders, and a set $\mathcal{N} = \mathcal{N}^{\Delta'}$ of M_N narrow cylinders.

Since each WWN-vertex $v \in V_{\Delta'}(\rho)$ (where $V_{\Delta'}(\rho)$ is the subset of all vertices of the union that lie inside Δ' for which ρ is a good direction) involves a wide cylinder, it must lie on the boundary of the sandwich region enclosed between the ρ-upper and ρ-lower envelopes of the strips in $\mathcal{T}^+_{\sigma'}(\rho)$ and $\mathcal{T}^-_{\sigma'}(\rho)$, respectively (recall that these sets are related only to the wide cylinders). Moreover, v is obtained as the intersection of an edge of the sandwich region and a narrow cylinder in Δ'.

We now apply a second decomposition step, in which we partition Δ' into sub-prisms Δ'' (clipped to within Δ'), so that each of them contains a relatively small number of both narrow and wide cylinders, as follows.

We draw two samples $R^+(\rho) \subseteq \mathcal{T}^+_{\sigma'}(\rho)$, $R^-(\rho) \subseteq \mathcal{T}^-_{\sigma'}(\rho)$ (satisfying the cutting property) of $O(r \log r)$ strips each, for some sufficiently large constant parameter r, and construct the overlay $\mathcal{M}_\rho$ of the respective minimization diagrams of the ρ-upper and ρ-lower envelopes of the strips in $R^+(\rho)$ and $R^-(\rho)$; these diagrams are obtained on a plane $H_{\rho^\perp}$ perpendicular to the ρ-direction. By the results of [2, 22], the overall complexity of $\mathcal{M}_\rho$ is $O(r^{2+\varepsilon})$, for any $\varepsilon > 0$. We next refine the overlay $\mathcal{M}_\rho$ by projecting the narrow cylinders in Δ' onto the plane $H_{\rho^\perp}$, taking a sample R (satisfying the cutting property) of $O(r \log r)$ of their bounding lines, and then overlaying $\mathcal{M}_\rho$ with the lines in R; we thus draw three samples in total. Let $\mathcal{M}^R_\rho$ be the refined decomposition. We first show:

Lemma 2.5 *The overall number of vertices in $\mathcal{M}^R_\rho$ is $O(r^{2+\varepsilon})$, for any $\varepsilon > 0$.*

Proof The overall complexity of the arrangement $\mathcal{A}(R)$ of the lines in R is $O(r^2 \log^2 r)$, and, as noted above, the overall complexity of $\mathcal{M}_\rho$ is $O(r^{2+\varepsilon})$, for any $\varepsilon > 0$. We next show that the overall number of edge-crossings between $\mathcal{M}_\rho$ and $\mathcal{A}(R)$ is $O(r \log r \lambda_6(r \log r)) = O(r^2 \operatorname{polylog} r)$, where $\lambda_s(q)$ is the maximal length of Davenport–Schinzel sequences of order s on q symbols (see [28]). This will assert the bound in the lemma.

Indeed, let $\mathcal{H}_R$ be the set of the $O(r \log r)$ planes containing the lines in R and parallel to the ρ-direction. Each plane $H \in \mathcal{H}_R$ intersects each of the strips in $R^+(\rho)$, $R^-(\rho)$ in an elliptic arc (clipped to within Δ'). Let Γ^+_H be the set of these ρ-upper arcs, and let Γ^-_H be the set of these ρ-lower arcs. Since H is parallel to the ρ-direction, it must meet (i) any edge e on the boundary of the ρ-upper envelope of the strips of $R^+(\rho)$ in a vertex that appears along the boundary of the upper envelope of the arcs in Γ^+_H (in the coordinate frame of H), (ii) any edge e' on the boundary of the ρ-lower envelope of the strips of $R^-(\rho)$ in a vertex that appears along the boundary of the lower envelope of the arcs in Γ^-_H (in the coordinate frame of H), and (iii) any edge e'' on the boundary of the sandwich region, enclosed between the ρ-upper envelope of the strips in $R^+(\rho)$ and the ρ-lower envelope of the strips in $R^-(\rho)$, in a vertex that appears along the boundary of the sandwich region enclosed between the upper envelope of the arcs in Γ^+_H and the lower envelope of the arcs in Γ^-_H (in the coordinate frame of H). Each of these vertices is projected on the plane $H_{\rho^\perp}$, perpendicular to the ρ-direction, to an edge-crossing between $\mathcal{M}_\rho$ and $\mathcal{A}(R)$.

We next use the following property. Let Q, Q' be two respective collections of m arcs each, where each pair of arcs intersect in at most s points (for some integer parameter $s > 0$). Then the number of vertices that appear in the overlay of the minimization diagrams of the respective upper and lower envelopes of these two collections is $O(\lambda_{s+2}(m))$. In fact, this is also the bound on the complexity of the envelope of each individual family, as well as the complexity of the sandwich region enclosed between these two envelopes; see, e.g., [28] for details. We thus conclude that since each pair of arcs in $\Gamma^+_H \cup \Gamma^-_H$ intersect in at most four points (as they are elliptic arcs), the overall number of the aforementioned vertices is $O(\lambda_6(r \log r))$. Summing this bound over all planes $H \in \mathcal{H}_R$ yields the asserted bound. □

We next triangulate each of the cells of $\mathcal{M}^R_\rho$ (in the coordinate frame of $H_{\rho^\perp}$) and lift each of these simplices in the ρ-direction, thereby obtaining a collection Ξ of

$O(r^{2+\varepsilon})$ triangular prism-subcells. We collect all these prisms that lie in the sandwich region enclosed between the ρ-upper envelope of the strips in $R^+(\rho)$ and the ρ-lower envelope of the strips in $R^-(\rho)$. Each of these prisms is bounded in the ρ^-- and ρ^+-directions (either by the boundary of Δ' or) by strips $\tau^+ \in R^+(\rho)$ appearing on the ρ-upper envelope and $\tau^- \in R^-(\rho)$ appearing on the ρ-lower envelope, respectively. We thus obtain a decomposition of the above sandwich region (within Δ') into $O(r^{2+\varepsilon})$ triangular prisms.

We observe that each cylinder K that is wide in Δ' continues to be wide in each clipped prism-subcell Δ'', since the width of Δ'' cannot exceed that of Δ' (or of Δ). By the cutting property, each prism-subcell Δ'' of the resulting decomposition is intersected by at most $O(M_W/r)$ strips in $\mathcal{T}^+_{\sigma'}(\rho)$ and by at most $O(M_W/r)$ strips in $\mathcal{T}^-_{\sigma'}(\rho)$, with a constant of proportionality that depends linearly on κ (recall that κ is a sufficiently large constant that represents the number of canonical strips resulting in the partition of a single cylinder boundary).

We next bound the number of wide cylinders within a prism-subcell Δ''. The strips in $\mathcal{T}^+_{\sigma'}(\rho) \cup \mathcal{T}^-_{\sigma'}(\rho)$ are said to be *good*, and all the remaining strips (of the wide cylinders within Δ') are said to be *bad*. Clearly, the number of wide cylinders K in Δ'' that have at least one good strip that meets Δ'' is at most M_W/r, since we can charge K to that good strip. The number of the remaining (wide) cylinders K in Δ'' cannot be bounded in the above manner, as the bad strips are not part of our sample, and thus the cutting properties do not apply for them. Nevertheless, we can ignore such cylinders K in the subsequent analysis of Δ''. This does not violate the bound on the number of (original) WWN-vertices v in Δ'', since each such vertex is obtained on the boundary of the sandwich region enclosed between the ρ-upper and ρ-lower envelopes of the good strips, and thus v does not lie on $\partial K \cap \Delta''$. We thus continue the subsequent analysis within Δ'' with at most M_W/r wide cylinders.

Each cylinder K that is narrow in Δ' either becomes wide in Δ'' or continues to be narrow in that cell. Let $\mathcal{L}_\rho$ be the set of the ρ-*silhouette* lines of the narrow cylinders within Δ' (that is, the locus of all the ρ-tangency points on the boundary of the cylinder). Applying similar arguments as above, a narrow cylinder K within a prism-subcell Δ'' of the decomposition must have a ρ-silhouette line ℓ, whose projection on $H_{\rho^\perp}$ meets the projection of Δ'' on that plane. We thus charge the crossing of K and Δ'' to that of the respective projections of ℓ and Δ'', and conclude that the number of narrow cylinders within Δ'' is at most M_N/r, by the cutting property. We have thus shown:

Lemma 2.6 *The number of wide cylinders within a subcell Δ'' of Δ' that contribute WWN-vertices to the union is at most M_W/r, and the number of narrow cylinders in that subcell is at most M_N/r.*

We bound the number of WWN-vertices of the union by applying a recursive decomposition scheme, where in each step we construct $O(r^{2+\varepsilon})$ prism-cells Δ'', using $(1/r)$-cutting of the good strips (taken from the current set of wide cylinders that we process in Δ') and the boundary projections of the narrow cylinders, as above, where each subproblem consists of at most M_W/r wide cylinders and at most M_N/r narrow cylinders. In each step we dispose immediately of all the *new* WWW-vertices in Δ'' (that is, vertices that were WWN in the parent cell Δ' and have just became WWW) and continue to process in recursion the remaining WWN-vertices. Thus cylinders that were narrow in the parent cell of Δ'' and became wide in Δ'' need not be processed in any further recursive step, the only wide cylinders that we process are those of the original problem. The recursion bottoms out when either $M_W \le c$, or $M_N \le c$, for some absolute constant $c \ge 3$. We then bound by brute-force the number of the (remaining) WWN-vertices and thus obtain an overall bound of $O(M_W^2 M_N)$, which is $O(M_W^2 + M_N)$ on the number of these vertices.

The number of new WWW-vertices in Δ'' is $O((M_N + M_W)M_W^{1+\delta})$, for any $\delta > 0$. Indeed, the number of new wide cylinders within Δ'' is at most M_N, and the number of the old wide cylinders (cylinders that were wide in the parent cell of Δ'') is at most M_W. Assume first that $M_N \le M_W$. Following the considerations given in the WWW-case, we conclude that the number of new WWW-vertices in Δ'' is $O(M_W^{2+\delta})$, for any $\delta > 0$. If $M_N > M_W$, we partition the set of the narrow cylinders into $\lceil \frac{M_N}{M_W} \rceil$ (roughly) equal subsets, each of which containing at most M_W elements. Each new WWW-vertex v involves a pair of old wide cylinders and a new wide cylinder within Δ''. We thus apply the nearly-quadratic bound in Δ'' on the set of the old wide cylinders and each of the subsets of new wide cylinders separately, and then sum these bounds over all subsets, thereby obtaining an overall bound of $O(M_N M_W^{1+\delta})$, for any $\delta > 0$ (see also [3, 16] for similar considerations). The bound now follows. Since r is constant, we obtain that the overall number of new WWW-vertices, over all prism-subcells Δ'', is also $O((M_N + M_W)M_W^{1+\delta})$.

Let $U_1(M_N, M_W)$ denote the maximum number of WWN-vertices that appear on the boundary of the union at a recursive step involving up to M_N narrow cylinders and M_W wide cylinders. Then U_1 satisfies the following recurrence:

$$U_1(M_N, M_W) \le \begin{cases} O((M_N + M_W)M_W^{1+\delta}) + O(r^{2+\varepsilon})U_1(\frac{M_N}{r}, \frac{M_W}{r}) \\ \quad \text{if } \min\{M_N, M_W\} > c, \\ O(M_N + M_W^2) \quad \text{if } \min\{M_N, M_W\} \le c, \end{cases}$$

where $\varepsilon > 0$, $\delta > 0$ are arbitrary, $c \ge 3$ is an appropriate constant, and where the constant of proportionality in the first expression depends on (ε, δ, κ, and on) r. It is straightforward to verify (see the arguments given earlier in this section and [16, 17]) that the solution of this recurrence is $U_1(M_N, M_W) = O((M_N + M_W)M_W^{1+\varepsilon})$, for any $\varepsilon > 0$ (slightly larger than δ in the nonrecursive term and the ε in the recursive term, but still arbitrarily close to 0), with a constant of proportionality that depends on ε and on κ.

Summing this bound over all $\Delta' = \Delta \cap \sigma'$ and over all good directions $\rho \in \mathcal{Z}$, we obtain that the overall number of WWN-vertices in Δ is $O((M_N + M_W)M_W^{1+\varepsilon})$, for any $\varepsilon > 0$, with a constant of proportionality that depends on ε and κ.

Remark The recursive mechanism is such that we sample the strips that belong to the wide cylinders and not the actual wide cylinders, since we have to take only those strips that have vertices v of the union, for which ρ is a good direction, and then recursively construct out of them the sandwich region that they induce. The role of the wide cylinders is in bounding the number of new WWW-vertices created in a subcell Δ'' of the current cell Δ'. We thus recurse with the wide cylinders but construct the cutting with their good strips.

The number of WNN-vertices of the union As in the WWN-case, we fix a good direction $\rho \in \mathcal{Z}$ and a subslab σ', and bound the number of WNN-vertices in each subcell $\Delta' = \Delta \cap \sigma'$ separately. We then sum these bounds over all these subcells and over all ρ. Here too, we observe that since each WNN-vertex $v \in V_\Delta(\rho)$ involves a wide cylinder, it must lie on the boundary of the sandwich region enclosed between the ρ-upper and ρ-lower envelopes of the strips in $\mathcal{T}^+_{\sigma'}(\rho)$ and $\mathcal{T}^-_{\sigma'}(\rho)$, respectively.

Using similar notation as in the WWN-case, we apply a similar decomposition, where we recursively construct $O(r^{2+\varepsilon})$ cells Δ'', each of which meets at most M_W/r wide cylinders and at most M_N/r narrow cylinders. At each recursive step, we dispose immediately of all the *new* WWW- and WWN-vertices (that is, vertices v that were WNN-vertices at the parent cell of Δ'' and have become either WWW- or WWN-vertices at Δ'') and continue to process in recursion the (remaining) WNN-vertices. Applying similar considerations as above, the overall number of new WWW- and WWN-vertices in the subcells Δ'', over all Δ'', is $O((M_N+M_W)M_N^{1+\delta})$, for any $\delta > 0$. At the bottom of the recursion (when either $M_W \le c$ or $M_N \le c$, for some absolute constant $c \ge 3$), we bound by brute-force the number of the (remaining) WNN-vertices, obtaining an overall bound of $O(M_W M_N^2)$, which is $O(M_W + M_N^2)$, on the number of these vertices. Letting $U_2(M_N, M_W)$ denote the maximum number of WNN-vertices that appear on the boundary of the union at a recursive step involving up to M_N narrow cylinders and M_W wide cylinders, we have:

$$U_2(M_N, M_W) \le \begin{cases} O((M_N + M_W)M_N^{1+\delta}) + O(r^{2+\varepsilon})U_2(\frac{M_N}{r}, \frac{M_W}{r}) \\ \quad \text{if } \min\{M_N, M_W\} > c, \\ O(M_N^2 + M_W) \quad \text{if } \min\{M_N, M_W\} \le c, \end{cases}$$

where, as above, $\varepsilon > 0$, $\delta > 0$ are arbitrary, $c \ge 3$ is an appropriate constant, and where the constant of proportionality in the first expression depends on (ε, δ, κ, and on) r. Here too, it is easy to verify that the solution of this recurrence is $U_2(M_N, M_W) = O((M_N + M_W)M_N^{1+\varepsilon})$, for any $\varepsilon > 0$ (slightly larger than δ in the nonrecursive term and the ε in the recursive term, but still arbitrarily close to 0), with a constant of proportionality that depends on ε and on κ.

Remarks (1) The WWN- and WNN-cases, albeit being similar, cannot be handled at the same recursive step and need to be analyzed separately. When a WNN-vertex v becomes WWN in a subcell Δ'', we need to consider the envelopes with respect to a (possibly, new) good direction ρ', not necessarily equal to ρ. These envelopes correspond to the old and the *new* wide cylinders in Δ'', and we need to continue with the recursive decomposition described for the WWN-case, according to the direction ρ', with the set of these wide cylinders and the narrow ones that have survived when passing from Δ' to Δ''.

(2) The only part of the analysis in which we explicitly use the geometric structure of the cylinders is in the WWW-case. The analysis for the WWN- and WNN-cases

only uses the property that the vertices lie on the boundary of the sandwich region enclosed between two envelopes of the good strips.

2.4 The Case of Congruent Cylinders

We now study the case where all the cylinders have equal radii and show that a simple specialization of our approach leads to a bound of $O(n^{2+\varepsilon})$, for any $\varepsilon > 0$, on the complexity of their union.

We only need to make the following simple observation. We apply the decomposition described in Sect. 2.2 and observe that, since all cylinders are of equal radii, *all* of them are either wide or narrow in a fixed cell Δ of the decomposition. Thus all vertices of the union inside Δ are either of type WWW or of type NNN. We thus construct this decomposition recursively, where in each step we dispose of all the new WWW-vertices within each subcell Δ, and continue to bound the number of NNN-vertices recursively. Here we do not need to apply the second decomposition step described for the WWN- and WNN-cases.

We note that these arguments can easily be extended to the case of nearly congruent cylinders, that is, when the radii of the cylinders are different but vary between some constant $\alpha < 1$ and 1. In this case, a cell Δ that contains both narrow and wide cylinders implies that the smallest radius r of a cylinder that meets Δ satisfies $r > \alpha w/2$, where w is the width of Δ. That is, a narrow cylinder cannot be "too narrow" in this case. Following the notation of Sect. 2.3, we elaborate Lemma 2.4 by refining the decomposition of the slab σ and adding at least $\frac{1}{\alpha \sin^2(\pi/\kappa)}$ planes parallel to H, thereby obtaining subslabs σ' of smaller width. The analysis then proceeds verbatim as in the original problem.

3 An Extension to the Case of Cigars

We now extend the analysis to the following case. We are given a set $\mathcal{S} = \{s_1, \ldots, s_n\}$ of n line-segments in 3-space and a set $\mathcal{B} = \{B_1, \ldots, B_n\}$ of n balls of arbitrary radii. The set $\mathcal{C} = \{s_1 \oplus B_1, \ldots, s_n \oplus B_n\}$ consisting of the *Minkowski sums* of $s_i \in S$ and $B_i \in B$ (that is, $s_i \oplus B_i = \{x + y \mid x \in s_i, y \in B_i\}$), for $i = 1, \ldots, n$, is the input set to our problem. The elements of $\mathcal{C}$ are also referred to as *cigars* in this case. We show:

Theorem 3.1 *The combinatorial complexity of the union of n cigars of arbitrary radii in 3-space is $O(n^{2+\varepsilon})$, for any $\varepsilon > 0$, where the constant of proportionality depends on ε. The bound is almost tight in the worst case.*

We devote this section to the proof of Theorem 3.1.

We first assume that the cigars in $\mathcal{C}$ are in general position. As discussed in [3], this excludes the cases where a pair of the segments in $\mathcal{S}$ are parallel or intersect, two cigars in $\mathcal{C}$ are tangent to each other, a curve of intersection of the boundaries of any two cigars is tangent to a third one, a triple intersection of the boundaries of the cigars lie on a circle separating the cylindrical and spherical portions of one of the cigars, and four boundaries meet at a point.

As a result, each vertex v of the union $\mathcal{U} = \bigcup_{C \in \mathcal{C}} C$ is the intersection point of either (i) three cylinder boundaries, (ii) a pair of cylinder boundaries and one spherical boundary, (iii) a cylinder boundary and a pair of spherical boundaries, or (iv) three spherical boundaries. We refer to each of these vertices as being of type *CCC*, *CCS*, *CSS*, or *SSS*, respectively. Vertices that are obtained by the intersection of the circle separating the cylindrical and spherical portions of one cigar, and the boundary of another cigar are discarded from further consideration, as the overall number of such vertices in the entire arrangement is $O(n^2)$. In the analysis of [3], these four cases are handled separately; however, using our new approach, we show that all these cases can be analyzed simultaneously. Note that in the SSS-case, each vertex v of the union lies on the boundary of the union of three distinct spheres, and thus the problem is reduced in this case to bounding the complexity of the union of $2n$ balls in 3-space, which is known to be $O(n^2)$ (see, e.g., [28]); however, this case is subsumed by our analysis.

We now apply a similar decomposition scheme to that of Sect. 2.2. That is, we project all the cigars in $\mathcal{C}$ onto the xy-plane and construct a $(1/r_0)$-cutting for the boundaries of these projections, for some sufficiently large constant parameter r_0, thereby obtaining a set of $O({r_0}^2 \log^2 r_0)$ pseudo-triangles, which we lift into vertical prisms.

We next use a similar definition of wideness as in the case of cylinders. That is, a cigar $C \in \mathcal{C}$ that intersects a prism-cell Δ is wide in Δ if its radius r is at least $w/2$ (where w is the width of Δ), and narrow otherwise. Here too, it is easy to verify that the silhouette of a narrow cigar in a cell Δ must intersect that cell.

We apply a similar recursive mechanism as in Sect. 2.3. The major difference in the analysis with respect to the case of cylinders is in bounding the number of WWW-vertices. We use here similar ideas to those presented in [3] for bounding the complexity of cigars of equal radii. Let κ be the same constant as in Sect. 2.3. We partition each of the cylinder boundaries of the cigars C into κ canonical strips as before, and we cover (both) hemispherical portions of C by $O(\kappa^2)$ spherical caps, each of opening angle at most π/κ, so that no point lies in more than a constant number of caps. We define a good direction for a cap in a similar manner as for strips (see conditions (i)–(ii) in the WWW-case in Sect. 2.3). The set of bad directions for such a spherical cap τ is again the union $B_1 \cup B_2$, where B_1 is defined as in the case of cylinders, and B_2 is defined as the spherical band consisting of all points at spherical distance at most $2\pi/\kappa$ from a great circle B_τ on $\mathbb{S}^2$, where B_τ corresponds to the plane parallel to the tangent plane of the cap τ at its center.

Following the reasonings in Sect. 2.3 and the notation there, each vertex v of the union has at least one good direction ρ, taken from a set $\mathcal{Z}$ of $O(1)$ directions. It is now easy to verify that the analog to Lemma 2.4 in this extended case continues to hold as well. That is, using the notation of Lemma 2.4, where now each of τ_1, τ_2, τ_3 is either a canonical strip or a spherical cap bounding a cigar K_1, K_2, K_3, respectively, when we enter into the cigar K_i, $i = 1, 2, 3$, through a point $u \in \tau_i \cap \sigma'$ in the direction parallel to ρ, we reach $\partial\sigma'$ before leaving K_i, $i = 1, 2, 3$.

Let us fix a prism-cell Δ of the decomposition and a good direction ρ, and let σ, σ', $V_\Delta(\rho)$, $\mathcal{T}^+_{\sigma'}(\rho)$, $\mathcal{T}^-_{\sigma'}(\rho)$ be as in Sect. 2.3. Following this notation, let $\mathcal{S}^+_{\sigma'}(\rho)$ be the set of the ρ-upper spherical caps that intersect the subslab σ' and contribute at least one vertex to $V_\Delta(\rho)$. We define $\mathcal{S}^-_{\sigma'}(\rho)$ in an analogous manner. It now follows that each vertex $v \in V_\Delta(\rho)$ must appear on the boundary of the sandwich region enclosed between the ρ-upper envelope of the strips and caps in $\mathcal{T}^+_{\sigma'}(\rho) \cup \mathcal{S}^+_{\sigma'}(\rho)$ and the ρ-lower envelope of the strips and caps in $\mathcal{T}^-_{\sigma'}(\rho) \cup \mathcal{S}^-_{\sigma'}(\rho)$. Repeating the analysis for all good directions ρ, we obtain a nearly-quadratic bound on the overall number of WWW-vertices of the union in a prism-cell Δ.

The remaining steps of the analysis follow almost verbatim from the analysis for the case of cylinders. In this case we need to replace $\mathcal{T}^+_{\sigma'}(\rho)$ with $\mathcal{T}^+_{\sigma'}(\rho) \cup \mathcal{S}^+_{\sigma'}(\rho)$, and $\mathcal{T}^-_{\sigma'}(\rho)$ with $\mathcal{T}^-_{\sigma'}(\rho) \cup \mathcal{S}^-_{\sigma'}(\rho)$. We leave the relatively easy details to the reader. This concludes the proof of Theorem 3.1.

4 Concluding Remarks and Open Problems

The only part of the analysis in which we have explicitly used the geometric structure of the cylinders is in the WWW-case. The analysis in the WWN- and WNN-cases only uses the property that the vertices lie on the boundary of the sandwich region enclosed between two envelopes of the good strips, where the major step of the analysis in these cases relies on the decomposition that we apply. Thus our machinery can be extended to any set of bodies in 3-space, for which there is a natural definition for wideness (reps., narrowness), from which one can derive similar properties to those of Lemma 2.4. One such problem concerns the union of cones. In this case we say that a cone is wide with respect to a fixed cell Δ if the silhouette of the cone does not meet Δ. Nevertheless, it is not known if Lemma 2.4 can be extended to this scenario, in order to show a nearly-quadratic upper bound on the number of WWW-vertices of the union within a cell Δ.

Another related problem is the case of *kreplach*, where in this setting, we modify the definitions from Sect. 3, so that now $\mathcal{S}$ consists of n pairwise-disjoint triangles. This problem was addressed by Agarwal and Sharir [3], who presented an upper bound of $O(n^{2+\varepsilon})$, for any $\varepsilon > 0$, on the complexity of the union where the balls in the Minkowski sum have the same radii. We note, however, that for arbitrary radii, the complexity is $\Omega(n^3)$ in the worst case. Indeed, consider a set of n disjoint (degenerate) triangles in 3-space that consists of three (roughly) equal-size classes. The first class consists of lines drawn in the plane $z = 0$ and parallel to the x-axis. The distance between each pair of consecutive lines is slightly smaller than 1. The second class is a copy of the first one, rotated by 90 degrees (so now these lines are parallel to the y-axis), and placed on the plane $z = 1$. The third class is a set of planes parallel to the xy-plane, placed within equal distances in the range $0 < z < 1$. We next construct the Minkowski sum of each of the lines (in either of the first two classes) with a ball of unit radius and obtain two respective classes of infinite cylinders; see Fig. 4(a). We leave the elements at the third class intact. Each such plane h intersects each of the cylinders from the first class in a horizontal strip (in the coordinate frame of h) and each of the cylinders at the second class in a vertical strip; see Fig. 4(b). It is easy to verify that the intersection points of these strips appear on the boundary of the union, and thus each plane h contributes $\Omega(n^2)$ vertices to the union, for a total of $\Omega(n^3)$ vertices. As is easily seen, this lower bound construction continues to hold when we

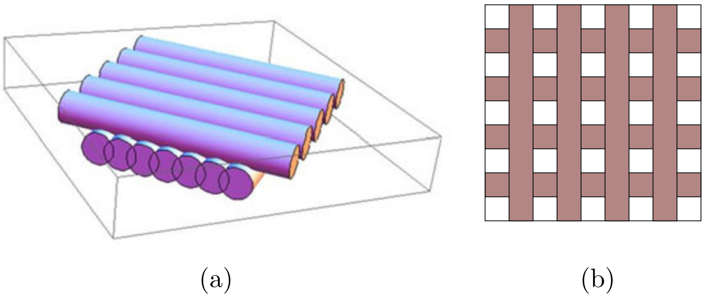

Fig. 4 (**a**) The Minkowski sum of a unit ball and each of the lines at the first two classes. (**b**) The cross section of these cylinders with a plane parallel to the xy-plane, lying between $z = 0$ and $z = 1$

replace each plane h with a sufficiently large triangle T (parallel to the xy-plane) and then construct the Minkowski sum of T and a ball of infinitesimally small radius.

Acknowledgements The author wishes to thank Micha Sharir for helpful discussions and his useful comments, and, in particular, for his help in writing this paper. The author also wishes to thank Boris Aronov and Pankaj Agarwal for useful discussions concerning the problem and, in particular, for Boris' useful comments about this paper. In addition, the author wishes to thank an anonymous referee for very useful comments regarding the presentation of this paper.

References

1. Agarwal, P.K., Pach, J., Sharir, M.: State of the union, of geometric objects: a review. In: Proc. Joint Summer Research Conference on Discrete and Computational Geometry—Twenty Years Later. Contemp. Math., vol. 452, pp. 9–48. AMS, Providence (2008)
2. Agarwal, P.K., Schwarzkopf, O., Sharir, M.: The overlay of lower envelopes and its applications. Discrete Comput. Geom. **15**(1), 1–13 (1996)
3. Agarwal, P.K., Sharir, M.: Pipes, cigars, and kreplach: The union of Minkowski sums in three dimensions. Discrete Comput. Geom. **24**, 645–685 (2000)
4. Aronov, B., Efrat, A., Koltun, V., Sharir, M.: On the union of κ-round objects in three and four dimensions. Discrete Comput. Geom. **36**, 511–526 (2006)
5. Aronov, B., Sharir, M.: On translational motion planning of a convex polyhedron in 3-space. SIAM J. Comput. **26**, 1785–1803 (1997)
6. Aronov, B., Sharir, M., Tagansky, B.: The union of convex polyhedra in three dimensions. SIAM J. Comput. **26**, 1670–1688 (1997)
7. Boissonnat, J.D., Sharir, M., Tagansky, B., Yvinec, M.: Voronoi diagrams in higher dimensions under certain polyhedral distance functions. Discrete Comput. Geom. **19**, 485–519 (1998)
8. Clarkson, K.L.: New applications of random sampling in computational geometry. Discrete Comput. Geom. **2**, 195–222 (1987)
9. Clarkson, K.L., Shor, P.W.: Applications of random sampling in computational geometry, II. Discrete Comput. Geom. **4**, 387–421 (1989)
10. de Berg, M.: Improved bounds on the union complexity of fat objects. Discrete Comput. Geom. **40**(1), 127–140 (2008)
11. de Berg, M.: Better bounds on the union complexity of locally fat objects. In: Proc. 26th Annu. Sympos. Comput. Geom., pp. 39–47 (2010)
12. Efrat, A.: The complexity of the union of (α, β)-covered objects. SIAM J. Comput. **34**, 775–787 (2005)
13. Efrat, A., Katz, M.: On the union of α-curved objects. Comput. Geom. Theory Appl. **14**, 241–254 (1999)

14. Efrat, A., Sharir, M.: On the complexity of the union of fat objects in the plane. Discrete Comput. Geom. **23**, 171–189 (2000)
15. Ezra, E., Aronov, B., Sharir, M.: Improved bound for the union of fat triangles. In: Proc. 22nd Annu. ACM–SIAM Sympos. Discr. Alg. (to appear)
16. Ezra, E., Sharir, M.: Almost tight bound for the union of fat tetrahedra in three dimensions. J. ACM **57**(1), 2 (2009)
17. Ezra, E., Sharir, M.: Counting and representing intersections among triangles in three dimensions. Comput. Geom. Theory Appl. **32**, 196–215 (2005)
18. Halperin, D.: On the complexity of a single cell in certain arrangements of surfaces related to motion planning. Discrete Comput. Geom. **11**, 1–33 (1994)
19. Halperin, D., Sharir, M.: Almost tight upper bounds for the single cell and zone problems in three dimensions. Discrete Comput. Geom. **14**(4), 385–410 (1995)
20. Haussler, D., Welzl, E.: ε-nets and simplex range queries. Discrete Comput. Geom. **2**, 127–151 (1987)
21. Kedem, K., Livne, R., Pach, J., Sharir, M.: On the union of Jordan regions and collision-free translational motion amidst polygonal obstacles. Discrete Comput. Geom. **1**, 59–71 (1986)
22. Koltun, V., Sharir, M.: The partition technique for overlays of envelopes. SIAM J. Comput. **32**(4), 841–863 (2003)
23. Matoušek, J., Miller, N., Pach, J., Sharir, M., Sifrony, S., Welzl, E.: Fat triangles determine linearly many holes. In: Proc. 32nd Annu. IEEE Sympos. Found. Comput. Sci., pp. 49–58 (1991)
24. Matoušek, J., Pach, J., Sharir, M., Sifrony, S., Welzl, E.: Fat triangles determine linearly many holes. SIAM J. Comput. **23**, 154–169 (1994)
25. Pach, J.: On the complexity of the union of geometric objects. Lect. Notes Comput. Sci. **2098**, 292–307 (2001)
26. Pach, J., Safruti, I., Sharir, M.: The union of congruent cubes in three dimensions. Discrete Comput. Geom. **30**, 133–160 (2003)
27. Pach, J., Tardos, G.: On the boundary complexity of the union of fat triangles. SIAM J. Comput. **31**, 1745–1760 (2002)
28. Sharir, M., Agarwal, P.K.: Davenport–Schinzel Sequences and Their Geometric Applications. Cambridge University Press, Cambridge (1995)