Lines Avoiding Balls in Three Dimensions Revisited

Natan Rubin

Received: 6 January 2010 / Revised: 3 January 2012 / Accepted: 27 January 2012 /

Published online: 23 February 2012

© Springer Science+Business Media, LLC 2012

Abstract Let $\mathcal{B}$ be a collection of n arbitrary balls in $\mathbb{R}^3$. We establish an almost-tight upper bound of $O(n^{3+\varepsilon})$, for any $\varepsilon > 0$, on the complexity of the space $\mathcal{F}(\mathcal{B})$ of all the lines that avoid all the members of $\mathcal{B}$. In particular, we prove that the balls of $\mathcal{B}$ admit $O(n^{3+\varepsilon})$ free isolated tangents, for any $\varepsilon > 0$. This generalizes the result of Agarwal et al. (Discrete Comput. Geom. 34:231–250, 2005), who established this bound only for *congruent* balls, and solves an open problem posed in that paper. Our bound almost meets the recent lower bound of $\Omega(n^3)$ of Glisse and Lazard (Proc. 26th Annu. Symp. Comput. Geom., pp. 48–57, 2010).

Our approach is constructive and yields an algorithm that computes the discrete representation of the boundary of $\mathcal{F}(\mathcal{B})$ in $O(n^{3+\varepsilon})$ time, for any $\varepsilon > 0$.

Keywords Lines in space · Tangents to spheres · Geometric arrangements · Combinatorial complexity · Free space · Union of simply-shaped bodies

1 Introduction

Let $\mathcal{B}$ be a collection of n arbitrary balls in 3-space. We show that the combinatorial complexity of the space $\mathcal{F}(\mathcal{B})$ of all lines that *miss* all the balls of $\mathcal{B}$, also referred to as *free* lines, is $O(n^{3+\varepsilon})$, for any $\varepsilon > 0$. See below for a more formal statement of the result. This result is related to several topics in computational and combinatorial geometry, including ray shooting and visibility in space, motion planning, the

N. Rubin (✉)

The Blavatnik School of Computer Science, Tel Aviv University, Tel Aviv 69978, Israel
e-mail: rubinnat@post.tau.ac.il

Present address:

N. Rubin

Department of Mathematics and Computer Science, Freie Universität Berlin, Takustr. 9 14195, Berlin, Germany

theory of arrangements, and the theory of transversals. Our result extends the result established by Agarwal et al. [1] for the case of unit (or congruent) balls, and solves an open problem posed in their paper. See [1] for a more detailed discussion of the problem.

1.1 Isolated Free Lines

A line in $\mathbb{R}^3$ is *isolated* (with respect to $\mathcal{B}$) if it is tangent to one or more balls in $\mathcal{B}$ and cannot be continuously moved without missing or penetrating into one of its tangent balls. Lines in $\mathbb{R}^3$ have *four* degrees of freedom, and they are naturally represented by points in the four-dimensional projective Grassmanian manifold; see [8] and [14, Sect. 7]. However, for the purpose of combinatorial analysis, we may assume that they are represented as points in Euclidean 4-space. This can be done in such a way that the combinatorial complexity of $\mathcal{F}(\mathcal{B})$ does not decrease; see [1] for similar analysis. Each ball $B \in \mathcal{B}$ defines an algebraic surface σ_B in $\mathbb{R}^4$ which is the locus of all (points representing) lines tangent to B . Since any line on the boundary of $\mathcal{F}(\mathcal{B})$ is necessarily tangent to some $B \in \mathcal{B}$, the boundary of $\mathcal{F}(\mathcal{B})$ is contained in the union of the surfaces σ_B , for $B \in \mathcal{B}$.

The combinatorial complexity of $\mathcal{F}(\mathcal{B})$ is defined as the number of faces of $\mathcal{A}$, of all dimensions, contained in $\partial\mathcal{F}(\mathcal{B})$. For simplicity of exposition, we will assume *general position* of the balls in $\mathcal{B}$, meaning that (i) the set of lines tangent to any k balls in $\mathcal{B}$, for $1 \leq k \leq 4$, corresponds to a $(4 - k)$ -dimensional surface in the parametric space $\mathbb{R}^4$, (ii) no line is tangent to any five distinct balls in $\mathcal{B}$, (iii) each plane is tangent to at most three balls in $\mathcal{B}$, and (iv) each point is contained in the boundaries of at most three balls in $\mathcal{B}$. This involves no loss of generality because we can apply a sufficiently small random perturbation to the balls, putting them in general position, without losing more than an (expected) constant fraction of the faces of $\partial\mathcal{F}(\mathcal{B})$; see Sect. 7.3.1 of [16] for similar arguments.

The lower-dimensional faces of $\partial\mathcal{F}(\mathcal{B})$ (of dimension 0, 1 or 2) represent lines tangent to two or more balls of $\mathcal{B}$. Since we are in four (parametric) dimensions, it is easy to see that the number of such faces is at most proportional to n^3 plus the number of boundary *vertices*, each of which is a common intersection of four of the surfaces σ_B , and represents a line which is tangent to the four corresponding balls (and misses all the other balls). Thus, any boundary vertex of $\mathcal{F}(\mathcal{B})$ corresponds to an isolated tangent line which misses the rest of the balls. Most of this paper is therefore devoted to proving a nearly cubic bound on the number of isolated free tangent lines.

1.2 Transversals and Free Lines

The tangency surface σ_B of each ball $B \in \mathcal{B}$ encloses the set σ_B^{stab} of all lines which stab (intersect) B . Thus, the space of all free lines corresponds to the complement of $\bigcup_{B \in \mathcal{B}} \sigma_B^{\text{stab}}$, and each isolated free tangent line corresponds to a boundary vertex of $\bigcup_{B \in \mathcal{B}} \sigma_B^{\text{stab}}$. In contrast, the *intersection* $\bigcap_{B \in \mathcal{B}} \sigma_B^{\text{stab}}$ corresponds to the *transversal space* $\mathcal{T}(\mathcal{B})$ of all line transversals of $\mathcal{B}$, which is, in a sense, antithetical to $\mathcal{F}(\mathcal{B})$. The complexity of $\mathcal{T}(\mathcal{B})$ has been extensively studied in the past two decades; see [17] for a comprehensive survey. Specifically, the above 4-dimensional representation, with a

careful choice of parameters, allows us to represent $\mathcal{T}(\mathcal{B})$ as a certain *sandwich region* in $\mathbb{R}^4$, enclosed between an upper envelope of a collection of surfaces and a lower envelope of another collection (where these surfaces are portions of the tangency surfaces σ_B); see [10, 16]. Thus, the nearly cubic bound on the complexity of such sandwich regions, as established in [11], immediately leads to a nearly cubic bound on the complexity of $\mathcal{T}(\mathcal{B})$ for any collection $\mathcal{B}$ of simply shaped convex sets (formally, of *constant description complexity*; see [11]) in $\mathbb{R}^3$. Unfortunately, unlike the transversal space $\mathcal{T}(\mathcal{B})$, the space $\mathcal{F}(\mathcal{B})$ of free lines cannot be conveniently represented as a sandwich region in $\mathbb{R}^4$, which makes the general bound of [11] inapplicable. For example, if $\mathcal{B}$ consists of n triangles, the complexity of $\mathcal{F}(\mathcal{B})$ can be as large as $\Theta(n^4)$ [15].

1.3 Our Results: Overview and Organization

As stated above, assuming general position of $\mathcal{B}$, it suffices to bound the number of isolated free tangent lines ℓ which are tangent to exactly four balls of $\mathcal{B}$, say, B_0, B_1, B_2, B_3 . Following the approach of [10], we *charge* ℓ to one of its tangent balls, say B_0 , and derive a bound of $O(n^{2+\varepsilon})$, for any $\varepsilon > 0$, on the number of isolated free tangent lines that are charged to a fixed ball $B_0 \in \mathcal{B}$. The somewhat involved charging rule depends on the geometric layout of the four tangent balls; this rule and its analysis are spelled out in detail in Sect. 2. This immediately implies the asserted bound of $O(n^{3+\varepsilon})$, for any $\varepsilon > 0$, on the number of free isolated tangent lines.

We thus reduce our study to the space of lines tangent to some fixed ball $B_0 \in \mathcal{B}$ and avoiding the interiors of all balls in $\mathcal{B} \setminus \{B_0\}$. The advantage of this reduction is that the restricted space $\mathcal{L} = \mathcal{L}(B_0)$ of lines tangent to B_0 has only three degrees of freedom, and corresponds to a three-dimensional manifold in $\mathbb{R}^4$ (again, ignoring the issue of projectivity; see [9, 10] for similar arguments). As in the general case, the space $\mathcal{F}(\mathcal{B} \setminus \{B_0\}) \cap \mathcal{L}$ of free lines within $\mathcal{L}$ can also be described as the complement, within $\mathcal{L}$, of $\bigcup_{B \in \mathcal{B} \setminus \{B_0\}} (\sigma_B^{\text{stab}} \cap \mathcal{L})$. Clearly, each isolated free tangent line corresponds to a vertex of the three-dimensional set $\mathcal{F}(\mathcal{B} \setminus \{B_0\}) \cap \mathcal{L}$. The task at hand, for a fixed B_0 , is therefore to bound the complexity of the union of the regions $\sigma_B^{\text{stab}} \cap \mathcal{L}$, for $B \in \mathcal{B} \setminus \{B_0\}$. A recent work of Ezra [5] deals with a similar problem, and shows that the complexity of the union of n arbitrary cylinders in $\mathbb{R}^3$ is $O(n^{2+\varepsilon})$, for any $\varepsilon > 0$. As it turns out, we can adapt the analysis technique of [5] and apply it to our setup (even though the geometric structures of these problems are quite different). The adaptation requires some care, because of the more involved geometry in our setup, but the overall resulting scheme is similar (and even somewhere simplified).

Informally, the analysis in [5] counts vertices of the union of the input cylinders. The analysis is recursive, and each sub-problem that it faces is defined within some portion τ of 3-space. Each vertex encountered in such a sub-problem is incident to the boundaries of three cylinders, each of which is classified as being either *wide* or *narrow*. The analysis bounds directly the number of WWW-vertices (incident to three wide cylinders) by representing them as vertices of some *sandwich region*, enclosed between the upper envelope of some collection of bivariate surfaces (portions of the cylinder boundaries) and the lower envelope of another such collection, within an

appropriate coordinate frame. The number of vertices of such a sandwich region is nearly quadratic [4, 11]. All the other kinds of vertex are bounded recursively; see [5] for details.

We use a similar scheme for each ball B_0 , working in the corresponding 3-dimension parametric space $\mathcal{L} = \mathcal{L}(B_0)$. That is, we want to represent isolated free tangent lines in $\mathcal{L}$ as vertices of appropriate sandwich regions. Similar to [5], this is tricky because these lines are vertices of the *union* of the “stabbing regions” $\sigma_B^{\text{stab}} \cap \mathcal{L}$, which, by itself, does not form a sandwich region in general. To overcome this main technical issue, several preparations have to be taken: (i) We need to define an appropriate parametrization of $\mathcal{L}$ and a good coordinate frame, within which the sandwich regions can be formed. (ii) We need to define the charging of lines ℓ to tangent balls B_0 so as to ensure the existence of such a good frame for ℓ within $\mathcal{L}(B_0)$, within which ℓ can be represented as a sandwich vertex. (iii) We need to set a recursive scheme, similar to the one in [5], so that lines in $\mathcal{L}(B_0)$ tangent to three other “wide” balls (see Sect. 3 for the precise definition) do indeed show up as vertices of sandwich regions, and so that the other lines can be handled recursively. All these steps are technically involved, and are spelled out in detail in Sects. 2 and 3 of the paper.

In Sect. 2 we describe a convenient way to represent lines tangent to the chosen ball B_0 , and establish several geometric facts, which are essential for adapting the method from [5]. In particular, we formally describe how to charge an isolated free tangent line to at least one of the balls in $\mathcal{B}$ it is tangent to, and how does this charging relate to the rest of the analysis. In Sect. 3 we derive a nearly quadratic bound on the number of isolated free tangent lines charged to any fixed ball B_0 , which implies our main result of a nearly cubic bound on the complexity of $\mathcal{F}(\mathcal{B})$. In Sect. 4 we describe an algorithm that computes $\mathcal{F}(\mathcal{B})$ in nearly cubic time, and give some concluding remarks and directions for future research.

2 Charging Free Lines to a Fixed Tangent Ball

As mentioned and detailed in the introduction, we assume general position of $\mathcal{B}$, which, in particular, implies that each isolated free tangent line is tangent to exactly four balls in $\mathcal{B}$. Let B_0 be a fixed ball in $\mathcal{B}$.

2.1 Lines Tangent to a Fixed Ball

First, we introduce several notations. Let $\mathcal{L} = \mathcal{L}(B_0)$ denote the space of all lines that are tangent to B_0 . Lines in $\mathcal{L}$ have *three* degrees of freedom, and we parametrize them as follows. The representation is with respect to a fixed choice of a *reference direction* $\vec{v} \in \mathbb{S}^2$; we will later choose different reference directions, and thereby obtain different parameterizations of $\mathcal{L}$. We first consider the *standard* representation, which uses the standard xyz -coordinate frame, in which $\vec{v}$ is the positive z -direction. We then represent each *oriented* line $\ell \in \mathcal{L}$ by the triple $(\theta(\ell), \kappa(\ell), \varphi(\ell))$, where $(\theta(\ell), \varphi(\ell)) \in \mathbb{S}^2$ is the representation in spherical coordinates of the orientation¹

¹That is, the positive z -direction is the zenith, θ gives the azimuth of ℓ with respect to the x -axis, and φ gives the elevation of ℓ .

Fig. 1 Representing $\ell \in \mathcal{L}$ by $(\theta(\ell), \kappa(\ell), \varphi(\ell))$. Three views: ℓ -parallel projection (left), xy -projection (middle), and inside the plane $h(\theta(\ell), \kappa(\ell))$ (right)

of ℓ , and $\kappa(\ell) \in (0, 2\pi)$ represents the point $\ell \cap B_0$ on the great circle of B_0 orthogonal to the direction $(\theta(\ell), \varphi(\ell))$. Here we set $\kappa(\ell) = 0$ for the line touching the great circle at its highest point (in the $\vec{v}$ -direction), and let it increase in the counterclockwise direction, with respect to the orientation of ℓ . The $\vec{v}$ -parallel (vertical) plane through the line ℓ with $\kappa(\ell) = 0$ passes through the center of B_0 . Note that each line $\ell \in \mathcal{L}$ with a fixed $\theta(\ell) = \theta \in (0, 2\pi)$ lies in one of the $\vec{v}$ -parallel planes making an angle θ with the x -axis in our coordinate frame.²

It is important to note that all lines with a fixed $(\theta(\ell), \kappa(\ell)) = (\theta, \kappa)$ lie in a single $\vec{v}$ -parallel plane, $h(\theta, \kappa)$, so that $\varphi(\ell) \in (0, \pi)$ describes the angle that ℓ makes with the $\vec{v}$ -parallel axis inside $h(\theta, \kappa)$. As φ varies from 0 to π , the line ℓ turns around the upper semicircle, if $-\pi/2 \leq \kappa \leq \pi/2$, or the lower semicircle, if $\pi/2 \leq \kappa \leq 3\pi/2$, of the disk $h(\theta, \kappa) \cap B_0$. See Fig. 1(right) for an illustration.

2.2 θ -span

Let $B \in \mathcal{B} \setminus \{B_0\}$ be a ball such that the xy -projections of B_0 and B are disjoint. Let I_B^+ and I_B^- be the pair of antipodal arcs on $\mathbb{S}^1$ such that $I_B^+ \cup I_B^-$ is the set of all $\theta \in \mathbb{S}^1$ for which there is at least one line $\ell \in \mathcal{L}$, with $\theta(\ell) = \theta$, which stabs B . Define the θ -span of B to be the angular extent θ_B of I_B^+ (or of I_B^-). Clearly, θ_B is equal to the angle between the two vertical planes that are common tangents to B_0 and B and separate the two balls. See Fig. 2 for an illustration. If the xy -projections of B_0 and B intersect, we either set θ_B to be π , when B is not contained in B_0 , or, otherwise, ignore B (within the analysis pertaining to the fixed ball B_0), since no $\ell \in \mathcal{L}$ can be tangent to B .

2.3 Motivation

The notion of the θ -span lies at the heart of our recursion scheme, detailed in Sect. 3. Recall from the Introduction that our goal is to bound the complexity of (the complement of) the union of the stabbing regions σ_B^{stab} , for $B \in \mathcal{B} \setminus \{B_0\}$. If the lines of $\mathcal{L}$ are represented in the standard $(\theta, \kappa, \varphi)$ -frame of $\mathcal{L}$ (which is induced by the original z -axis), the θ -span θ_B of B measures the extent of the θ -projection $I_B^+ \cup I_B^-$ of the stabbing region σ_B^{stab} of B (within $\mathcal{L}$).

²We ignore all lines ℓ that are parallel to the $\vec{v}$ -direction, since their values $\theta(\ell)$ are not well-defined. Assuming that the choice of $\vec{v}$ is generic, no isolated free lines are lost.

Fig. 2 The θ -span of B equals the angle made by two vertical planes that are tangent to B_0 and B and separate B_0 from B

Fig. 3 Isolated free lines $\ell \in \mathcal{L}$ correspond to boundary vertices of $\bigcup_{B \in \mathcal{B} \setminus \{B_0\}} \sigma_B^{\text{stab}}$. By restricting $\theta(\ell) \in \Delta$ we focus on such vertices (*isolated free lines*) within the corresponding θ -vertical slab. Here B is wide

At each step of our main recursion we analyze the number of isolated free lines³ $\ell \in \mathcal{L}$ whose θ -coordinates $\theta(\ell)$ are restricted to some fixed interval Δ_0 . The respective vertices of $\bigcup_{B \in \mathcal{B} \setminus \{B_0\}} \sigma_B^{\text{stab}}$ lie within the θ -vertical slab that projects to Δ_0 . We partition Δ_0 into sub-intervals Δ , each yielding an even more restricted instance of the overall problem; see Fig. 3. For each sub-interval $\Delta \subset \Delta_0$ we separately bound the number of isolated free lines ℓ with $\theta(\ell) \in \Delta$. A ball B is classified as wide (resp., narrow) in Δ if its θ -span is larger (resp., smaller) than the extent of Δ (and one of I_B^+ , I_B^- intersects Δ). According to this definition, a ball can be narrow in at most four sub-intervals Δ (two such intervals for each of the arcs I_B^- , I_B^+). This facilitates efficient recursive treatment of the isolated free lines that are tangent to three narrow balls (within Δ).

For each of the remaining intervals Δ , either the stabbing region σ_B^{stab} totally avoids their θ -vertical slabs, or B is wide in Δ . In the former case, the ball B can be simply discarded, as there are no isolated free lines tangent to B with $\theta(\ell) \in \Delta$. In the latter case, the wideness of B implies a certain “local” monotonicity property of the region σ_B^{stab} (within Δ), which we discuss later in this section. This allows to bound the number of isolated free lines that are tangent to at least one such wide ball more directly. For instance, if an isolated free line ℓ (with $\theta(\ell) \in \Delta$) is tangent to *three* wide balls (in addition to B_0), it can be represented as a vertex of a suitably defined sandwich region within Δ . In the end, *every* isolated free line is represented as a sandwich vertex in some recursive sub-problem. To enable this representation, we need a new coordinate frame $(\theta', \kappa', \varphi')$, which is determined by a new reference direction $\vec{v}$ (different from the standard $\vec{z}$). In addition, this requires a more careful selection of B_0 among the four balls that are tangent to the chosen isolated free line ℓ (we shall refer to this process as *charging* ℓ to B_0).

³To be pedantic, an additional restriction will be imposed on these lines. We shall assign (or *charge*) every isolated free line to one of its four tangent balls. In our analysis we shall only consider the lines that are charged to the ball B_0 which defines $\mathcal{L}$ (rather than all isolated free lines that are tangent to B_0).

Fig. 4 Choosing a good reference direction $\vec{v}$. All lines $\ell \in \mathcal{L}$ that have same (θ', κ') -coordinates as ℓ , lie in the plane $H(\ell, \vec{v})$ (left). As the line $\ell(\gamma)$ rotates in the plane $H(\ell, \vec{v})$, its direction moves along the $\vec{v}$ -orthogonal great circle C_H (right)

2.4 Choosing a Good Reference Direction

In the remainder of this paper, $(\theta, \kappa, \varphi)$ denotes the *standard* coordinate frame of $\mathcal{L}$, i.e., the one corresponding, as above, to the reference direction $(0, 0, 1)$ and to the x - and y -axes complementing the frame.

In Sect. 3 we construct a collection of $O(1)$ reference directions. Each such direction $\vec{v} \in \mathbb{S}^2$ defines a *new* coordinate frame $(\theta', \kappa', \varphi')$ in which *some* isolated free lines $\ell \in \mathcal{L}$ become boundary vertices $(\theta'(\ell), \kappa'(\ell), \varphi'(\ell))$ of certain *sandwich regions* (see [4, 11] and the introduction). Any such region is enclosed between an upper envelope of one collection of $\theta'\kappa'$ -monotone portions of the tangency surfaces of the balls of $\mathcal{B} \setminus \{B_0\}$, restricted to the 3-dimensional space $\mathcal{L}$, and a lower envelope of another such collection. In Sect. 3 we give a formal definition of these sandwich regions and argue that their total complexity is bounded by $O(n^{2+\varepsilon})$, for any $\varepsilon > 0$. It hence suffices to make sure that each isolated free line $\ell \in \mathcal{L}$ can be represented as a sandwich vertex for at least one reference direction, which is said to be *good* for $\ell \in \mathcal{L}$. To ensure this property, we charge each isolated free tangent line in $\mathcal{F}(\mathcal{B})$ to (at least) one of the four balls it is tangent to, and establish the existence of good reference directions only for those lines charged to B_0 .

More formally, fix some sufficiently small constant $\delta > 0$. Let $B \in \mathcal{B} \setminus \{B_0\}$ be a ball with θ -span θ_B , and let $\ell \in \mathcal{L}$ be a line tangent to B . Let $\vec{v} \in \mathbb{S}^2$ be a reference direction, and let $(\theta', \kappa', \varphi')$ be the coordinate frame of $\mathcal{L}$ which corresponds to $\vec{v}$.⁴ For each $-\varphi'(\ell) \leq \gamma \leq \pi - \varphi'(\ell)$ let $\ell(\gamma)$ denote the line whose $(\theta', \kappa', \varphi')$ -coordinates are $(\theta'(\ell), \kappa'(\ell), \varphi'(\ell) + \gamma)$. Observe that $\ell(0) = \ell$ and that all such lines $\ell(\gamma)$ lie in the $\vec{v}$ -parallel plane through ℓ ; let us denote this plane by $H(\ell, \vec{v})$. Keeping θ', κ' fixed and varying γ causes the line $\ell(\gamma)$ to rotate around the disk $B_0 \cap H(\ell, \vec{v})$. See Fig. 4. An appropriate choice of the direction of the rotation (either clockwise or counterclockwise, starting from $\ell(0) = \ell$) will cause $\ell(\gamma)$ to penetrate into the interior of B as soon as it starts moving away from ℓ . We keep rotating $\ell(\gamma)$ in this direction until either it stops penetrating B and becomes tangent to it again, or it becomes parallel to the $\vec{v}$ -axis, in which case $\varphi'(\ell) + \gamma$ becomes equal to either 0 or π . Letting $\tilde{\ell}$ denote the final position of $\ell(\gamma)$, we say that the direction $\vec{v}$ is *good* for the pair (ℓ, B) if either $|\theta(\ell) - \theta(\tilde{\ell})| > \delta\theta_B$ or $\varphi'(\tilde{\ell}) \in \{0, \pi\}$.

⁴The choice of the other two $\vec{v}$ -orthogonal axes of the frame is arbitrary; see the remark below.

Fig. 5 View of ℓ , $\tilde{\ell}$ and σ_B in the standard coordinate frame (left), and in the new one (right). The θ -orthogonal strip Γ_ℓ does not contain the second tangent $\tilde{\ell}$ (left). The free space within Γ_ℓ and near $(\theta'(\ell), \kappa'(\ell), \varphi'(\ell))$ lies entirely below σ_B (right)

Informally, this goodness means that, if we keep $(\theta', \kappa') = (\theta'(\ell), \kappa'(\ell))$ fixed and vary φ' into σ_B^+ (i.e., rotate ℓ so as to make it penetrate into B), then ℓ will remain nonfree for a “long stretch”, meaning that its (standard) θ -coordinate changes by at least $\delta\theta_B$ (or ℓ reaches some simpler terminal position with $\varphi' \in \{0, \pi\}$) before ℓ gets out of B again. This will provide the basis for making ℓ a sandwich vertex.

Remark Note that the other tangent, $\tilde{\ell}$, is fully determined by the choice of $\vec{v}$. Hence the choice of the two $\vec{v}$ -orthogonal coordinates which, together with $\vec{v}$, form the new coordinate system in $\mathbb{R}^3$, is immaterial for the definition of a good direction.

Good Reference Directions Imply $\theta'\kappa'$ -Monotonicity Let $\vec{v}$ be a good direction for the pair (ℓ, B) , as above. Let Γ_ℓ denote the θ -orthogonal strip of lines having their standard θ -coordinates within $(\theta(\ell) - \delta\theta_B, \theta(\ell) + \delta\theta_B)$; see Fig. 5(left). By the choice of $\vec{v}$ and working in the new coordinate frame $(\theta', \kappa', \varphi')$, the ball B is stabbed by all lines in Γ_ℓ and above $(\theta'(\ell), \kappa'(\ell), \varphi'(\ell))$ in the φ' -direction, or B is stabbed by all such lines below $(\theta'(\ell), \kappa'(\ell), \varphi'(\ell))$. That is, if we keep $(\theta', \kappa') = (\theta'(\ell), \kappa'(\ell))$ fixed and vary φ' from $\varphi'(\ell)$, say, upwards, then we reach the boundary of Γ_ℓ before exiting σ_B^{stab} (at $\tilde{\ell}$). See Fig. 5.

That is, a good reference direction for (ℓ, B) defines a coordinate frame $(\theta', \kappa', \varphi')$ in which the free space lies entirely above, or entirely below the surface σ_B near ℓ . See Fig. 5(right). More precisely, all this is true for the θ -vertical strip of width $2\delta\theta_B$ around ℓ .

Definition We say that the reference direction $\vec{v}$ is *good* for an isolated free line $\ell \in \mathcal{L}$ tangent to B_0 and to some three other balls B_i , $1 \leq i \leq 3$, if $\vec{v}$ is good for each of the pairs (ℓ, B_i) , $1 \leq i \leq 3$, with B_0 as the “base” ball.

Informally, the existence of a good direction $\vec{v}$ for a line $\ell \in \mathcal{L}$ can be exploited in the following way. Let Δ be a θ -interval (say, generated by the main recursion scheme, as reviewed above and detailed in Sect. 3) such that $\theta(\ell) \in \Delta$. Suppose ℓ is tangent to three balls $B_1, B_2, B_3 \in \mathcal{B} \setminus \{B_0\}$ so that⁵ $\delta\theta_{B_i} > |\Delta|$ (for each $1 \leq i \leq 3$). In the coordinate frame $(\theta', \kappa', \varphi')$ induced by $\vec{v}$, the free space over Δ lies

⁵Here $|\Delta|$ denotes the angular extent of Δ .

entirely above, or entirely below each tangency surface σ_{B_i} (near ℓ), for $1 \leq i \leq 3$. Hence, ℓ corresponds to a vertex of a sandwich region which is constructed from (appropriately clipped) $\theta'\kappa'$ -monotone portions of the surfaces σ_B , for $B \in \mathcal{B} \setminus \{B_0\}$. (For each choice of Δ and $\vec{v}$ we shall construct only one sandwich region.) See Sect. 3 for more details.

2.5 Charging Lines to Balls

Our strategy is to choose $O(1)$ canonical directions $\vec{v}$ within $\mathcal{L}(B_0)$, for each ball $B_0 \in \mathcal{B}$, and charge each isolated free tangent line $\ell \in \mathcal{F}(\mathcal{B})$ to one of the balls, B_0 , it is tangent to, so that we are guaranteed to have a canonical direction $\vec{v}$ which is good for (ℓ, B) , for each of the three other balls B tangent to ℓ . This is a somewhat intricate step, which proceeds as follows.

Definition For each free isolated line ℓ tangent to B_0 and to some three other balls $B_1, B_2, B_3 \in \mathcal{B}$, we charge ℓ to B_0 if the set of all planes h containing ℓ , such that the disk $B_0 \cap h$ has radius *smaller* than or equal to the radii of the disks $B_i \cap h$, for all $1 \leq i \leq 3$, contains a double wedge of angle at least $\pi/6$; see Fig. 6(left) for an illustration.

Lemma 2.1 *Each isolated free tangent line is charged to at least one of the four balls in $\mathcal{B}$ it is tangent to.*

Proof Parametrize the (oriented) planes containing ℓ by Π_θ , where $0 \leq \theta \leq 2\pi$ measures the angle between the normal to the plane and some fixed plane normal. Let B_i , $0 \leq i \leq 3$, denote the four balls tangent to ℓ . For each $0 \leq i \leq 3$ and $0 \leq \theta \leq 2\pi$ we define $r_i(\theta)$ to be the radius of the circle $B_i \cap \Pi_\theta$. See Fig. 6(right) for an illustration. We claim that, for any pair $i < j$, $r_i(\theta) = r_j(\theta)$ has two solutions, up to antipodality (clearly, if θ is a solution so is $\theta + \pi$). To see this, project ℓ , B_i and B_j onto some plane orthogonal to ℓ . The respective projections B_i^* , B_j^* of B_i , B_j are two discs whose boundaries meet at the point ℓ^* to which ℓ projects; see Fig. 6(right). Each plane Π_θ projects to a line Π_θ^* through ℓ^* , and $r_i(\theta)$ is half the chord $\Pi_\theta^* \cap B_i^*$, and similarly for B_j . Hence we have $r_i(\theta) = \frac{1}{2}R_i|\cos(\theta - \alpha_i)|$, where R_i is the radius of B_i and α_i is the angle at which Π_θ^* intersects B_i^* in a diameter, and similarly for $r_j(\theta)$. This implies that $r_i(\theta) = r_j(\theta)$ when $R_i \cos(\theta - \alpha_i) = R_j \cos(\theta - \alpha_j)$ or $R_i \cos(\theta - \alpha_i) = -R_j \cos(\theta - \alpha_j)$, and each of these equations has a unique solution in the range $(0, \pi)$ and an antipodal solution in $(\pi, 2\pi)$. This implies the claim.

Hence, in the range $(0, \pi)$, the four functions $r_i(\theta)$ behave as pseudo-parabolas, and therefore their minimization diagram consists of at most $\lambda_2(4) = 7$ arcs,⁶ delimited by 6 breakpoints; see [16]. Hence, over $(0, 2\pi)$ the diagram consists of at most 12 arcs, which appear in antipodal pairs. Hence, there must exist such a pair for which the angular extent of each arc in the pair is at least $2\pi/12 = \pi/6$. By definition, we can charge ℓ to the ball B_i which attains the lower envelope at this pair of arcs. $\square$

⁶In general, $\lambda_s(n)$ denotes the maximum length of a Davenport–Schinzel sequence of order s on n symbols. For the special case $s = 2$ we have $\lambda_2(n) = 2n - 1$; see [16].

Fig. 6 Charging an isolated free tangent line ℓ to a tangent ball B_0 ; the angle of the shaded double wedge is $\geq \pi/6$ (left); the ball B_1 is “responsible” for both lines delimiting the wedge. The projections B_i^* , B_j^* of balls B_i , B_j , for $0 \leq i < j \leq 3$, and the respective functions $r_i(\theta)$, $r_j(\theta)$ —the equation $r_i(\theta) = r_j(\theta)$ has two solutions, θ_1 and θ_2 , up to antipodality (right)

The following main theorem of this section shows that our choice of the charged ball B_0 , for each isolated free line ℓ , does imply the existence of a good reference direction for ℓ .

Theorem 2.2 *Let η be a sufficiently small positive constant and let $\ell \in \mathcal{L}$ be an isolated free line charged to B_0 , which is tangent to B_0 and to three other balls B_i , $1 \leq i \leq 3$, so that the center c_i of B_i lies at distance d_i from the center c_0 of B_0 , for $i = 1, 2, 3$. Suppose that for each $1 \leq i \leq 3$ the angle between c_0c_i and the z -axis is at least $\sin^{-1}\eta$. Then there is a spherical cap $C \subset \mathbb{S}^2$ of radius $1/100$, say, such that all directions $\vec{v} \in C$ are good reference directions for ℓ (with a constant δ that depends on η ; specifically, we can choose $\delta = \Theta(\eta^3)$).*

Clearly, the set of good directions also contains the spherical cap antipodal to C .

Proof We start with the set $\mathbb{S}^2$ and gradually reduce the set of directions under consideration, until we end up with a subset of $\mathbb{S}^2$ all of whose directions $\vec{v}$ are good for ℓ , which still contains a spherical cap of radius $1/100$. As before, we assume that B_0 is centered at the origin and has unit radius.

Since ℓ is charged to B_0 , Lemma 2.1 implies the existence of a double wedge W_0 of planes through ℓ , of angle $\geq \pi/6$ with the following property: For any plane $H \subset W_0$, the disk $B_0 \cap H$ has radius smaller than that of $B_i \cap H$, for all $1 \leq i \leq 3$. We restrict ourselves to the directions $\vec{v} \in \mathbb{S}^2$ for which the corresponding plane $H(\ell, \vec{v})$ is contained in W_0 . On the unit sphere $\mathbb{S}^2$, those directions form a spherical double wedge bounded by the two great circles parallel to the planes bounding W_0 ; see Fig. 7. Thus, restricting $\vec{v}$ to lie in this double wedge, we ensure that the radius of $B_0 \cap H(\ell, \vec{v})$ is smaller than the radii of the disks $B_i \cap H(\ell, \vec{v})$, for $1 \leq i \leq 3$.

We next ensure that for each $1 \leq i \leq 3$ the disk $B_i \cap H(\ell, \vec{v})$ has radius at least ηr_i , where r_i denotes the radius of B_i . Let H_i denote the plane that is tangent to B_i at the point $\ell \cap B_i$, and let C_i denote the H_i -parallel great circle on $\mathbb{S}^2$. Clearly, for each $0 \leq i \leq 3$ the disk $H(\ell, \vec{v}) \cap B_i$ has radius $r_i \sin \alpha$, where α is the dihedral angle between $H(\ell, \vec{v})$ and H_i . Note that α is greater than or equal to the angle between $\vec{v}$ and H_i . This follows from the easily established fact that the dihedral angle between the planes maximizes the angle between a vector on one plane and the other plane; see Fig. 8(left).

Fig. 7 The constructed set of good directions is contained in a spherical double wedge W bounded by the two great circles parallel to the planes bounding W_0

Fig. 8 The dihedral angle α between a pair of planes H_1, H_2 is greater than or equal to the angle β between a vector $\vec{v}$ in H_1 and H_2 (left). The case when $B_0 \cap H$ intersects $B_i \cap H$: The line $\ell(\gamma)$ must rotate by an angle greater than π (while intersecting B) before it becomes tangent to B_i again (right)

Thus, in order to ensure that the radius of $B_i \cap H(\ell, \vec{v})$ is greater than ηr_i it is sufficient to remove all directions $\vec{v}$ which lie in the strip of width $2 \sin^{-1} \eta = O(\eta)$ around C_i , for each $1 \leq i \leq 3$.

Finally, we exclude all directions $\vec{v}$ which make an angle of at most $\sin^{-1} \eta$ with the z -parallel plane $H^\perp = H(\ell, (0, 0, 1))$ that contains ℓ . All the removed directions lie in a strip of width $2 \sin^{-1} \eta$ around the $H^\perp$ -parallel great circle in $\mathbb{S}^2$. Each remaining direction $\vec{v}$ makes an angle of at least $\sin^{-1} \eta$ with $H^\perp$ (and, in particular, with the z -axis).

Clearly, by choosing η sufficiently small, the set of all unremoved reference directions $\vec{v} \in \mathbb{S}^2$ contains a spherical cap C of radius $1/100$. We next claim that any surviving reference direction (and in particular any direction $\vec{v} \in C$) is good for ℓ . Recall that this is equivalent to showing that $\vec{v}$ is good for each of the pairs (ℓ, B_i) , for $1 \leq i \leq 3$.

To establish the claim, fix the ball B_i , for some $1 \leq i \leq 3$, and a reference direction, $\vec{v} \in C$. Let $(\theta', \kappa', \varphi')$ be the coordinate frame that corresponds to $\vec{v}$, and let $H = H(\ell, \vec{v})$ denote the $\vec{v}$ -parallel plane through ℓ . Suppose, without loss of generality, that the rotated line $\ell(\gamma)$ penetrates the interior of B_i for arbitrarily small $\gamma > 0$. If $\ell(\gamma)$ stabs B_i for all $0 \leq \gamma \leq \pi - \varphi'(\ell)$ then $\vec{v}$ is good for (ℓ, B_i) by definition. Thus, we may assume that for some $0 \leq \gamma \leq \pi - \varphi'(\ell)$ the line $\ell(\gamma)$ becomes tangent to B_i again, and leaves it immediately afterwards. In particular, this, and the fact that $r(B_i \cap H) > r(B_0 \cap H)$, implies that $B_0 \cap H$ is disjoint from $B_i \cap H$. Otherwise, ℓ is a common outer tangent to $B_0 \cap H$ and $B_i \cap H$, and $\ell(\gamma)$ must rotate around B_0 by

⁷We use the inequality $\sin^{-1} \eta < 2\eta$, say, for η sufficiently small, and choose η in the statement of the theorem to satisfy this inequality.

Fig. 9 The lines ℓ and $\tilde{\ell} = \ell(\Delta_{\text{end}})$ (left). The θ -interval $(\theta(\ell), \theta(\tilde{\ell}))$ (right)

an angle greater than π , while intersecting B_i , before it becomes tangent to B_i again, so it must become $\vec{v}$ -parallel before leaving B_i ; see Fig. 8(right).

Assume then that $B_0 \cap H$ and $B_i \cap H$ are disjoint, and let γ_{end} be the smallest $0 \leq \gamma \leq \pi - \varphi'(\ell)$ such that $\ell(\gamma)$ misses the interior of B_i , and put $\tilde{\ell} = \ell(\gamma_{\text{end}})$. In the plane $H = H(\ell, \vec{v})$, the line $\tilde{\ell}$ is another common tangent to the disks $B_0 \cap H$ and $B_i \cap H$, and makes an angle of γ_{end} with ℓ ; see Fig. 9(left). Note that one of $\ell, \tilde{\ell}$ is an outer tangent and the other is a common inner tangent to these two disks.

We need to show that $|\theta(\tilde{\ell}) - \theta(\ell)| > \delta \theta_{B_i}$, for an appropriate constant $\delta > 0$ (which will depend on η). This is done in three steps. First, we show that $\gamma_{\text{end}} > \eta r_i / d_i$. Then we prove that θ_{B_i} is at most $O(r_i / (\eta d_i))$. Finally, we show that $|\theta(\tilde{\ell}) - \theta(\ell)|$ is $\Omega(\eta \gamma_{\text{end}})$. The combination of these three inequalities would then imply that $|\theta(\tilde{\ell}) - \theta(\ell)| > \delta \theta_{B_i}$, with $\delta = \Omega(\eta^3)$. We proceed to establish the three claims:

Claim 1: $\gamma_{\text{end}} > \eta r_i / d_i$.

Proof Refer to Fig. 9(left), depicting the cross-section of B_0 and B_i within $H = H(\ell, \vec{v})$. Denote the center of $B_0 \cap H$ by c_0^* and its radius by r_0^* . Similarly, denote the center of $B_i \cap H$ by c_i^* and its radius by r_i^* . We have $c_0^* c_i^* \leq d_i$ and $r_0^* \leq r_i^*$. We also have $r_i^* \geq \eta r_i$, by construction of $\vec{v}$. Draw through c_0^* lines parallel to the two tangents $\ell, \tilde{\ell}$. One of $\ell, \tilde{\ell}$ is a common inner tangent to $B_0 \cap H, B_i \cap H$, say it is $\tilde{\ell}$. Denote by γ_0 the angle between $c_0^* c_i^*$ and the line through c_0^* parallel to $\tilde{\ell}$; see Fig. 9(left). The angle between the two new lines (parallel to $\ell, \tilde{\ell}$) is equal to γ_{end} , and since $r_0^* \leq r_i^*$, it is greater than γ_0 . Hence⁸

$$\gamma_{\text{end}} \geq \gamma_0 = \sin^{-1} \frac{r_0^* + r_i^*}{|c_0^* c_i^*|} \geq \sin^{-1} \frac{r_0^* + r_i^*}{d_i} \geq \sin^{-1} \frac{\eta r_i}{d_i} > \frac{\eta r_i}{d_i}.$$

In this argument we assume that $B_0 \cap H$ and $B_i \cap H$ are disjoint, so $\tilde{\ell}$ exists, which is indeed what the current analysis has been assuming. $\square$

Claim 2: $\theta_{B_i} = O(r_i / (\eta d_i))$.

Proof Here we project B_0 and B_i on the xy -plane. Referring to Fig. 10(left), it is easily seen that the distance $\hat{d}_i$ between the centers of the projection disks of B_0 and

⁸The last inequality in the following derivation follows from the inequality $\sin x < x$ for $x > 0$.

Fig. 10 If the xy -projections of the two balls are disjoint then $\sin \theta_{B_i}/2 = (r_0 + r_i)/\hat{d}_i$ (left). Proving that B_0 has radius at most $14r_i$ (right)

B_i is at most d_i . On the other hand, the angle between c_0c_i and the z -axis is at least $\sin^{-1} \eta$, by assumption, so we have $\hat{d}_i \geq \eta d_i$. Moreover, the radii of the projection disks are clearly r_0 and r_i , respectively.

We claim that r_0 cannot be much bigger than r_i . Specifically, we show that $r_0 \leq 14r_i$. Indeed, recall that ℓ is charged to B_0 , so there exists a double wedge W_0 bounded by two planes passing through ℓ and forming a dihedral angle of at least $\pi/6$, such that, for every plane H contained in W_0 , $r(B_0 \cap H) \leq r(B_i \cap H)$. Project W_0 , B_0 , B_i and ℓ on some plane orthogonal to ℓ . In the projection, B_0 and B_i are mapped to two respective disks B_0^* , B_i^* , so that their boundaries meet at the point ℓ^* into which ℓ projects; see Fig. 10(right).

Suppose to the contrary that $r_0 > 14r_i$. Draw a line λ through ℓ^* which makes equal and openly disjoint chords within B_0^* and B_i^* . Draw another line μ by connecting ℓ^* to the other intersection point of ∂B_0^* and ∂B_i^* . Let τ_0 denote the tangent to B_0^* at ℓ^* . One of the double wedges, W_1 , delimited by λ and μ consists of lines which mark in B_0^* a smaller chord than the chord marked in B_i^* and W_1 has to contain τ_0 , because the “chord” it makes in B_0^* is of length 0; see Fig. 10 (right). The length of each of the two chords $\lambda \cap B_0^*$, $\mu \cap B_0^*$ is at most the diameter of B_i^* , namely at most $2r_i$. Therefore, the angles β_1, β_2 between these chords and the tangent τ_0 , which are half the respective central angles subtended by these chords at the center of B_0^* , are small. Specifically, we have $\sin \beta_1, \sin \beta_2 \leq r_i/r_0 < 1/14$. Hence the angular size of W_1 is $\beta_1 + \beta_2 < \pi/6$. However, W_1 , when lifted back to 3-space, is exactly the wedge W_0 with the property that $r(B_0 \cap H) < r(B_i \cap H)$ for every plane H in W_0 . By assumption, the dihedral angle of W_0 is at least $\pi/6$, a contradiction that shows that $r_0 \leq 14r_i$.

If the xy -projections of B_0 and B_i are disjoint, as depicted in Fig. 10(left), we have

$$\sin \frac{\theta_{B_i}}{2} = \frac{r_0 + r_i}{\hat{d}_i} \leq \frac{r_0 + r_i}{\eta d_i} \leq \frac{15r_i}{\eta d_i},$$

implying that $\theta_{B_i} = O(r_i/(\eta d_i))$.

If the xy -projections of B_0 and B_i intersect, then $\theta_{B_i} = \pi = \Theta(1)$ and $\hat{d}_i \leq r_i + r_0 \leq 15r_i$. The claim now follows since

$$\frac{r_i}{\eta d_i} \geq \frac{r_i}{\hat{d}_i} \geq \frac{r_i}{15r_i} \geq \frac{1}{15}. \quad \square$$

Claim 3: $|\theta(\tilde{\ell}) - \theta(\ell)| = \Omega(\eta \gamma_{\text{end}})$.

Proof Let α denote the dihedral angle between $H = H(\ell, \vec{v})$ and the xy -plane. By construction of $\vec{v}$, $\alpha < \pi/2 - \sin^{-1} \eta$. We have two lines $\ell, \tilde{\ell}$ in H forming an angle

γ_{end} between them, and we need to bound from below the angle between their xy -projections.

Without loss of generality, assume that H meets the xy -plane at the x -axis, and let v be the unit vector in H orthogonal to the x -axis. That is $v = (0, \cos \alpha, \sin \alpha)$. Let $e = (1, 0, 0)$ denote the unit vector in the x -direction. Let ω, ω' denote the signed angles that $\ell, \tilde{\ell}$ form with the x -axis (so $|\omega - \omega'| = \gamma_{\text{end}}$). Then the unit vectors in the directions of $\ell, \tilde{\ell}$ are, respectively,

$$w = e_1 \cos \omega + v \sin \omega = (\cos \omega, \sin \omega \cos \alpha, \sin \omega \sin \alpha)$$

and

$$w' = e_1 \cos \omega' + v \sin \omega' = (\cos \omega', \sin \omega' \cos \alpha, \sin \omega' \sin \alpha).$$

Hence, the unit vectors in the directions of the xy -projections of $\ell, \tilde{\ell}$ are, respectively,

$$\frac{(\cos \omega, \sin \omega \cos \alpha)}{\sqrt{\cos^2 \omega + \sin^2 \omega \cos^2 \alpha}} \quad \text{and} \quad \frac{(\cos \omega', \sin \omega' \cos \alpha)}{\sqrt{\cos^2 \omega' + \sin^2 \omega' \cos^2 \alpha}}.$$

Hence, using the vector product, we have (up to a sign change)

$$\begin{aligned} \sin(\theta(\ell) - \theta(\tilde{\ell})) &= \frac{(\cos \omega \sin \omega' - \cos \omega' \sin \omega) \cos \alpha}{\sqrt{\cos^2 \omega' + \sin^2 \omega' \cos^2 \alpha} \cdot \sqrt{\cos^2 \omega + \sin^2 \omega \cos^2 \alpha}} \\ &= \frac{\sin \gamma_{\text{end}} \cos \alpha}{\sqrt{\cos^2 \omega' + \sin^2 \omega' \cos^2 \alpha} \cdot \sqrt{\cos^2 \omega + \sin^2 \omega \cos^2 \alpha}}. \end{aligned}$$

We have $\cos \alpha > \cos(\pi/2 - \sin^{-1} \eta) = \eta$, and $\cos^2 \omega + \sin^2 \omega \cos^2 \alpha, \cos^2 \omega' + \sin^2 \omega' \cos^2 \alpha < 1$. Hence $\sin(\theta(\ell) - \theta(\tilde{\ell})) \geq \eta \sin \gamma_{\text{end}}$, from which the claim follows. $\square$

As noted, the three claims just proved complete the proof of Theorem 2.2. $\square$

Remark

- (1) The technical assumption in Theorem 2.2 about the angles between the vectors $c_0 c_i$ and the z -axis can be removed by using a constant number of different standard frames. Put differently, the definition of a good reference direction depends on the standard coordinate frame, and each isolated free tangent line ℓ may need to use a different standard frame (out of $O(1)$ possible ones) to deduce the existence of a large set of good directions for ℓ .
- (2) An additional pruning step makes each $\vec{v} \in C$ satisfy that the angle between $c_0 c_i$ and the $\vec{v}$ -direction is at least $\sin^{-1} \eta$, for each $1 \leq i \leq 3$. This is done by pruning all directions $\vec{v} \in \mathbb{S}^2$ that make angle at most $\sin^{-1} \eta$ with $c_0 c_i$, for $1 \leq i \leq 3$. Clearly, all of the pruned directions belong to the union of six spherical caps of radius $\sin^{-1} \eta$ on $\mathbb{S}^2$, so we can still assume that, for sufficiently small $\eta > 0$, the set of remaining directions contains a spherical cap of radius $1/100$. (The usefulness of this note will become clear when we prove Lemmas 3.5 and 3.7.)
- (3) In some steps of the analysis in Sect. 3 we will need to assume that $\vec{v}$ is a good reference direction for a free tangent line ℓ , in the strong sense that $\vec{v}$ is a direction surviving the pruning steps of our construction.

2.6 Counting Isolated Free Lines

In Sect. 3, we bound the number of lines charged to a fixed ball $B_0 \in \mathcal{B}$, which touches them. The proof proceeds by considering $O(1)$ candidate reference directions and bounding, for each such direction $\vec{v}$, the number of isolated free lines such that $\vec{v}$ is good for them. This analysis culminates in the following theorem.

Theorem 2.3 *Let $\mathcal{B}$ be a collection of n arbitrarily balls in $\mathbb{R}^3$ in general position. Then each ball $B \in \mathcal{B}$ is charged by at most $O(n^{2+\varepsilon})$ isolated free tangent lines, for any $\varepsilon > 0$.*

As an immediate corollary of Lemma 2.1 and Theorem 2.3, we obtain the following main result of the paper.

Theorem 2.4 *Let $\mathcal{B}$ be a collection of n arbitrary balls in $\mathbb{R}^3$ in general position. Then $\mathcal{B}$ admits $O(n^{3+\varepsilon})$ isolated free tangent lines, for any $\varepsilon > 0$.*

As mentioned, a recent construction of Glisse and Lazard [7] gives a set $\mathcal{B}$ of n balls (in general position) with $\Omega(n^3)$ isolated free tangent lines, thus making the bound in Theorem 2.4 almost tight in the worst case.

3 The Number of Lines Charged to a Fixed Ball

In this section we prove Theorem 2.3. Let $\eta > 0$ be a fixed constant as in Theorem 2.2. We restrict our task to bounding the number of isolated free lines tangent to B_0 and to some three other balls B_i , for $1 \leq i \leq 3$, which are charged to B_0 . We denote by d_{B_i} the distance of the center c_{B_i} of B_i from the center c_0 of B_0 , and assume that the angle between $c_0c_{B_i}$ and the z -axis is at least $\sin^{-1} \eta$, for each $1 \leq i \leq 3$. Assuming η is small enough, as remarked above, this restriction involves no loss of generality, since we can choose a constant number of directions for the z -axis such that any isolated free tangent line $\ell \in \mathcal{L}$ satisfies these inequalities for at least one of these directions. Therefore, we henceforth assume that the angle between $c_0c_{B_i}$ and the z -axis is at least $\sin^{-1} \eta$, for every $B \in \mathcal{B} \setminus \{B_0\}$.

In addition, we may assume, without loss of generality, that, for each $B \in \mathcal{B} \setminus \{B_0\}$, there is at least one isolated free line in $\mathcal{L}$ which is tangent to B and charged to B_0 .

Notation Recall from Sect. 2 that each ball $B \in \mathcal{B} \setminus \{B_0\}$ defines a pair of mutually antipodal circular arcs I_B^+ and I_B^- on $\mathbb{S}^1$ of angular extent θ_B each. Let $\Delta \subset \mathbb{S}^1$ be a circular arc of angular extent $\theta_\Delta < \pi$. We say that a ball $B \in \mathcal{B}$, having θ -extent θ_B , is *wide* (resp., *narrow*) in Δ if Δ intersects $I_B^+ \cup I_B^-$, and $\theta_\Delta < \theta_B$ (resp., $\theta_\Delta \geq \theta_B$); see Fig. 11.

Let $\ell \in \mathcal{L}$ be an isolated free tangent line with $\theta(\ell) \in \Delta$, which is tangent to B_0 and to some three other balls $B_i \in \mathcal{B} \setminus \{B_0\}$, $1 \leq i \leq 3$, and which is charged to B_0 . Clearly, for each $1 \leq i \leq 3$ the ball B_i is either wide or narrow in Δ . We classify ℓ as

Fig. 11 The circular arcs γ_1 and γ_3 are narrow with respect to $\Delta \subset \mathbb{S}^1$, while the arcs γ_2 and γ_4 are wide

a *free WWW-line* in Δ , if all three balls B_i , $1 \leq i \leq 3$, are wide in Δ . We classify ℓ as a *free WWN-line* (resp., *WNN-line*, *NNN-line*) in Δ , if exactly one (resp., two, three) balls among B_i , $1 \leq i \leq 3$, are narrow in Δ .

To establish the asserted nearly quadratic bound on the number of isolated free tangent lines that are charged to B_0 , we will analyze separately the number of isolated free tangents of each type. The overall analysis involves a three-level recursion scheme which resembles (and is actually simpler than) the one described in [5]; see also [6].

3.1 Overall Recursion

Let r be a sufficiently large constant (which will be determined later). In each step, we are given a circular arc $\Delta_0 \subset \mathbb{S}^1$ and a set $\mathcal{B}_N^{\Delta_0}$ of n narrow balls in Δ_0 , and our task is to bound the number of isolated free tangents in $\{B_0\} \cup \mathcal{B}_N^{\Delta_0}$ which (belong to $\mathcal{L}$ and) are charged to B_0 . (Here n denotes the size of the current sub-problem rather than that of the overall initial scenario.) For this, we decompose Δ_0 into at most $2r$ subarcs Δ , each containing at most n/r endpoints of the arcs I_B^+ , I_B^- , for $B \in \mathcal{B}_N^{\Delta_0}$, and each of extent smaller than π . At the first recursive step, Δ_0 is equal to $\mathbb{S}^1$, and $\mathcal{B}_N^{\Delta_0}$ is equal to $\mathcal{B} \setminus \{B_0\}$. (Here, and only here, the angular extent of Δ_0 is larger than π , but this will not affect the analysis.) As $\mathbb{S}^1$ is progressively cut into subarcs Δ , progressively more narrow balls become wide in these arcs, and the size of the subsets $\mathcal{B}_N^\Delta$ keeps decreasing. In fact, in a single recursive partition of an arc, a ball can remain narrow in at most two of the subarcs. During each step of recursion we immediately dispose of any *new* free WWW-, WWN-, and WNN- lines within each subarc Δ (which were free NNN-lines in the parent arc Δ_0), and continue to bound recursively the number of the remaining free NNN-lines. The recursion bottoms out when $n < r$. In this case the number of remaining isolated free tangent lines within the current arc Δ_0 can be crudely upper bounded by $O(r^4) = O(1)$.

We show below that the overall number of isolated free tangent lines of types WWW, WWN, and WNN, in a fixed subarc Δ_0 , is $O(n^{2+\varepsilon})$, for any $\varepsilon > 0$. Assume for the moment that we have these bounds. Let $F(\Delta_0, \mathcal{B}_N^{\Delta_0})$ denote the number of isolated free tangent lines $\ell \in \mathcal{L}$ which are tangent to B_0 and to some triple of balls in $\mathcal{B}_N^{\Delta_0}$ and which are charged to B_0 , and let $F(n)$ denote the maximum value of $F(\Delta_0, \mathcal{B}_N^{\Delta_0})$, over all possible choices of B_0 , a circular arc Δ_0 , and a collection $\mathcal{B}_N^{\Delta_0}$ of n narrow balls in Δ_0 . Then $F(n)$ satisfies the following recurrence:

$$F(n) \leq \begin{cases} 2rF(n/r) + O(n^{2+\varepsilon_0}), & n \geq r, \\ O(1), & n < r, \end{cases} \quad (1)$$

where $r > 0$ is a sufficiently large constant, $\varepsilon_0 > 0$ is an arbitrarily small constant, and the constants of proportionality in the non-recursive terms depend on r and ε_0 .

Clearly, for any $\varepsilon > 0$, choosing r sufficiently large and ε_0 sufficiently small (as functions of ε), yields the solution $F(n) = O(n^{2+\varepsilon})$, where the constant of proportionality depends on ε .

The remainder of this section is dedicated to deriving a nearly quadratic bound on the number of free WWW-, WWN-, and WNN-lines within a fixed subarc Δ of Δ_0 . By Theorem 2.2, there exists a set $\mathcal{D}$ of $O(1)$ directions on $\mathbb{S}^2$, so that for each free isolated tangent line ℓ in $\mathcal{L}$ (which is charged to B_0) there exists $\vec{v} \in \mathcal{D}$ which is a good reference direction for ℓ . (For example, we can choose a collection of $O(1)$ directions that hits every spherical cap of radius $1/100$. Recall that we may also assume (by applying the foregoing analysis to a constant number of choices of the z -direction) that the extra condition on the orientations of the vectors $c_0 c_B$ is satisfied for all balls $B \in \mathcal{B} \setminus \{B_0\}$.) Therefore, it is sufficient to fix a reference direction $\vec{v} \in \mathcal{D}$, and to bound the number of isolated free tangent lines of the above types, for which $\vec{v}$ is a good reference direction. Let $(\theta', \kappa', \varphi')$ denote the coordinate frame of $\mathcal{L}$ induced by $\vec{v}$, and let $(\theta, \kappa, \varphi)$ continue to denote the *standard* coordinate frame induced by the vertical direction $(0, 0, 1)$. Let δ be the constant from Theorem 2.2. We partition the fixed arc Δ into $\lceil 1/\delta \rceil$ subarcs $\Delta' \subset \Delta$, each of angular extent⁹ $\theta_{\Delta'} \leq \delta \theta_{\Delta}$. In what follows, we also fix a subarc Δ' , and consider only lines $\ell \in \mathcal{L}$ with $\theta(\ell) \in \Delta'$.

3.2 The Number of WWW-, WWN- and WNN-Lines: Overview

Informally, the analysis proceeds as follows. We first derive a nearly quadratic bound on the number of WWW-lines using a direct, non-recursive argument. Specifically, each WWW-line ℓ (having good direction $\vec{v}$) is tangent to three wide balls B_1, B_2, B_3 whose stabbing regions satisfy the desired “local” monotonicity property near ℓ (over Δ), as reviewed in Sect. 2. (This property follows from the goodness of $\vec{v}$ for ℓ and the wideness of B_1, B_2, B_3 in Δ .) This allows to represent ℓ as a vertex of an appropriately defined sandwich region within the *new* coordinate frame $(\theta', \kappa', \varphi')$ induced by $\vec{v}$. The general result of [4] provides a nearly quadratic upper bound on the number of such vertices.

Then we derive a recurrence for the number of WWN-lines, in which we use the bound on WWW-lines, and show that it also yields a nearly quadratic bound. Once again, we take advantage of the fact that each WWN-line with good direction $\vec{v}$ is tangent to two wide balls, whose stabbing regions are monotone near ℓ (over Δ). As in the case of WWW-lines, this requires working in the new coordinate frame $(\theta', \kappa', \varphi')$. Informally, this sub-recurrence causes the third ball to become wide in some sub-problem, and then the bound for WWW-lines can be used.

Finally, we derive another sub-recurrence for the number of WNN-lines, which is very similar to the first one and uses the bounds obtained in the two preceding steps. This recurrence too yields a nearly quadratic bound.

All these bounds are then plugged into the analysis given above, thereby obtaining an overall nearly quadratic bound for the number of isolated free lines that are charged to B_0 .

⁹Without loss of generality, we assume that $\delta < 1/2$, which ensures that each Δ' has θ -span at most π .

Fig. 12 The lines ℓ_0 and ℓ_1 , sharing the same $\theta'\kappa'$ -coordinates and the same plane $H = H(\ell_0, \vec{v}) = H(\ell_1, \vec{v})$, are, respectively, negative and positive tangents to B (left). The point $\theta(\ell_1)$ lies outside Δ' , implying that B is stabbed by all lines ℓ with $\theta'\kappa'\varphi'$ -coordinates $(\theta'_0, \kappa'_0, \xi)$, such that $\xi \leq \varphi'_0$ and $\theta(\ell) \in \Delta'$ (right)

3.3 The Number of Free WWW-Lines

Fix the arcs $\Delta' \subset \Delta \subset \Delta_0$ (and the reference direction $\vec{v}$), as above, and denote by $\mathcal{B}_W^\Delta$ and $\mathcal{B}_N^\Delta$ the subsets of $\mathcal{B}_N^{\Delta_0}$ consisting of balls which are, respectively, wide and narrow in Δ . Working in the $\theta'\kappa'\varphi'$ -coordinate frame of $\mathcal{L}$, we wish to represent each free WWW-line $\ell \in \mathcal{L}$ (which is charged to B_0 and satisfies $\theta(\ell) \in \Delta'$) as a vertex in a certain 3-dimensional *sandwich region* [16], below the *lower envelope* of some collection Γ^+ of at most n $\theta'\kappa'$ -monotone surfaces, and above the *upper envelope* of another collection Γ^- of at most n $\theta'\kappa'$ -monotone surfaces, where each surface consists of points representing common tangents to B_0 and to some other ball $B \in \mathcal{B}_W^\Delta$. Using the bounds in [4], this will yield the bound $O(n^{2+\varepsilon})$ on the number of free WWW-lines.

Definition A line $\ell \in \mathcal{L}$ is a *positive* (resp., *negative*) tangent to $B \in \mathcal{B} \setminus \{B_0\}$ if an infinitesimally small increase (resp., decrease) in its φ' -coordinate, while keeping the $\theta'\kappa'$ -coordinates fixed, causes ℓ to penetrate the interior of B . See Fig. 12(left).

For each $B \in \mathcal{B}_W^\Delta$ let σ_B^+ (resp., σ_B^-) denote the surface patch consisting of all points representing (in the $\theta'\kappa'\varphi'$ -frame) lines $\ell \in \mathcal{L}$ which are positive (resp., negative) tangents to B and satisfy the following additional conditions: (i) $\theta(\ell) \in \Delta'$, and (ii) $\vec{v}$ is a good reference direction for (ℓ, B) . See Fig. 12(left). As is easily verified, after appropriate re-parametrization, each surface patch σ_B^+, σ_B^- has *constant description complexity* (that is, it is defined as a Boolean combination of a constant number of polynomial equalities and inequalities, of constant maximum degree; see, e.g., [16]). The collection Γ^+ (resp., Γ^-) is the set of all (non-empty) surface patches of the form σ_B^+ (resp., σ_B^-), over $B \in \mathcal{B}_W^\Delta$.

The following lemma guarantees that for each ball $B \in \mathcal{B}_W^\Delta$ the surface patches $\sigma_B^+ \in \Gamma^+$ and $\sigma_B^- \in \Gamma^-$ defined above are $\theta'\kappa'$ -monotone and lie, respectively, above and below the set of all free lines $\ell \in \mathcal{L}$ with $\theta(\ell) \in \Delta'$.

Lemma 3.1 *Let B be a ball in $\mathcal{B}_W^\Delta$ and let $(\theta'_0, \kappa'_0, \varphi'_0)$ be a point on the surface patch $\sigma_B^+ \in \Gamma^+$ (resp., $\sigma_B^- \in \Gamma^-$). Then B is stabbed by all lines $\ell \in \mathcal{L}$ with $\theta'\kappa'\varphi'$ -coordinates $(\theta'_0, \kappa'_0, \xi)$, such that $\xi > \varphi'_0$ (resp., $\xi < \varphi'_0$) and $\theta(\ell) \in \Delta'$.*

Proof We establish the claim only for $(\theta'_0, \kappa'_0, \varphi'_0)$ lying on σ_B^+ ; the argument for σ_B^- is symmetric. Suppose B is missed by some line with coordinates $(\theta'_0, \kappa'_0, \xi)$, where $\xi > \varphi'_0$, and let φ'_1 be the infimum of all such ξ . Let ℓ_0 and ℓ_1 be the lines with respective coordinates $(\theta'_0, \kappa'_0, \varphi'_0)$ and $(\theta'_0, \kappa'_0, \varphi'_1)$. Then ℓ_0 is a positive tangent and ℓ_1 is a negative tangent to B . Since $\vec{v}$ is good for (ℓ_0, B) , we have $|\theta(\ell_0) - \theta(\ell_1)| \geq \delta\theta_B > \delta\theta_\Delta \geq \theta_{\Delta'}$. Thus, $\theta(\ell_1)$ lies outside Δ' .¹⁰

Note that the lines $\ell \in \mathcal{L}$ with coordinates $(\theta'_0, \kappa'_0, \xi)$ stab B for all $\varphi'_0 \leq \xi \leq \varphi'_1$. Now, referring to Fig. 12(right), we observe that as we rotate ℓ around B_0 by increasing $\varphi'(\ell) = \xi$ from φ'_0 to φ'_1 , the value $\theta(\ell)$ either monotonically increases or monotonically decreases.¹¹ Let ℓ be a line as in the lemma. Since $\theta(\ell_1)$ is outside Δ' , the condition that $\theta(\ell) \in \Delta'$ ensures that $\varphi'(\ell) \leq \varphi'_1 = \varphi'(\ell_1)$, implying that ℓ stabs B ; see Fig. 12(right). $\square$

We next argue that each free WWW-line $\ell \in \mathcal{L}$, such that ℓ is charged to B_0 , $\theta(\ell) \in \Delta'$, and $\vec{v}$ is a good direction for ℓ , corresponds to a vertex on the sandwich region below the lower envelope of Γ^+ and above the upper envelope of Γ^- . Indeed, any such line ℓ is tangent to B_0 , and to some three other balls $B_1, B_2, B_3 \in \mathcal{B}_W^\Delta$, such that $\vec{v}$ is good for (ℓ, B_i) , for each $1 \leq i \leq 3$. Therefore, $(\theta'(\ell), \kappa'(\ell), \varphi'(\ell))$ belongs to one of $\sigma_{B_i}^+, \sigma_{B_i}^-$ for each $1 \leq i \leq 3$. Moreover, by Lemma 3.1, $(\theta'(\ell), \kappa'(\ell), \varphi'(\ell))$ lies above all the other surfaces σ_B^- and below all the other surfaces σ_B^+ , for $B \in \mathcal{B}_W^\Delta$. Hence ℓ corresponds to a vertex of the sandwich region, as claimed.

It follows that the number of free WWW-tangents $\ell \in \mathcal{L}$ which are charged to B_0 , and satisfy the above conditions, is bounded by the number of boundary vertices of the above sandwich region. As already discussed, since each of Γ^+, Γ^- consists of $O(n)$ surface patches of constant description complexity, it follows from [4] (see also [11]) that the number of such vertices is $O(n^{2+\varepsilon})$, for any $\varepsilon > 0$. Thus, we have shown the following.

Theorem 3.2 *Let Δ be a circular arc on $\mathbb{S}^1$ and let $\mathcal{B}_W^\Delta$ be a collection of n wide balls in Δ . Then $\{B_0\} \cup \mathcal{B}_W^\Delta$ admits at most $O(n^{2+\varepsilon})$ free isolated tangent lines that are charged to B_0 , for any $\varepsilon > 0$.*

3.4 The Number of Free WWN-Lines

We continue to follow the setup assumed in the preceding analysis. That is, we have fixed arcs $\Delta' \subseteq \Delta \subseteq \Delta_0$ on $\mathbb{S}^1$, a fixed reference direction $\vec{v} \in \mathcal{D}$, and two sets $\mathcal{B}_W^\Delta, \mathcal{B}_N^\Delta$ of balls which are, respectively, wide and narrow in Δ . Put $m = |\mathcal{B}_W^\Delta|$ and $n = |\mathcal{B}_N^\Delta|$. We only consider free WWN-lines ℓ charged to B_0 and tangent to two balls in $\mathcal{B}_W^\Delta$ and to one ball in $\mathcal{B}_N^\Delta$, so that $\theta(\ell) \in \Delta'$, and $\vec{v}$ is a good reference direction for ℓ .

¹⁰The intersection between Δ' and $(\theta(\ell_0), \theta(\ell_1))$ is connected since $\theta_{\Delta'}, \theta_B < \pi$. (This condition has to be verified only at the top level of recursion; it holds trivially at deeper levels.)

¹¹This follows since the direction vector of ℓ moves along the $H(\ell, \vec{v})$ -parallel great circle on $\mathbb{S}^2$.

We continue working in the new¹² coordinate frame $(\theta', \kappa', \varphi')$ induced by $\vec{v}$. To bound the number of WWN-lines (under the above restrictions), we develop a secondary recursion scheme which decomposes the $\theta'\kappa'$ -domain into pseudo-trapezoidal regions. Each such region τ defines a (sub-)recursion instance which involves WWN-lines ℓ whose respective values $(\theta'(\ell), \kappa'(\ell))$ lie in τ . At each step of this secondary recursion the number of relevant¹³ balls from *each* set $\mathcal{B}_W^\Delta, \mathcal{B}_N^\Delta$ decreases by a factor of $\Omega(r/\log r)$. To achieve this, we define the decomposition of τ as an overlay of two decompositions. The first decomposition guarantees that the number of relevant balls from $\mathcal{B}_W^\Delta$ decreases by a factor of $\Omega(r/\log r)$. Its construction uses the previously defined $\theta'\kappa'$ -monotone surfaces of Γ^+, Γ^- .¹⁴ The second decomposition is very similar to the decomposition used by the overall recursion. It partitions the $\theta'\kappa'$ -domain into θ' -vertical strips (so the space $\mathcal{L}$ is partitioned into θ' -vertical slabs). It guarantees that the number of relevant balls from $\mathcal{B}_N^\Delta$ decreases at each step by a factor of $\Omega(r)$. Its definition is based on a suitably modified notion of wideness which involves the θ' -spans of these balls (rather than their θ -spans). See below for details.

For each $B \in \mathcal{B} \setminus \{B_0\}$, in analogy to the definition of the θ -span (which appears in Sect. 2), let $I_B^{'+}$ and $I_B'^{-}$ denote the pair of mutually antipodal circular arcs whose union forms the set of all θ' -values of lines $\ell \in \mathcal{L}$ that stab B . Accordingly, we define the θ' -span θ'_B of each $B \in \mathcal{B} \setminus \{B_0\}$ as the angular extent of each one of the arcs $I_B^{'+}, I_B'^{-}$. (As in the case of the θ -span, if the $\vec{v}$ -orthogonal projections of B_0 and B intersect each other, we either set θ'_B to be π , if B is not contained in B_0 , or, otherwise, ignore B , since no $\ell \in \mathcal{L}$ can be tangent to B . See Sect. 2 for more details.) Our analysis relies on the following technical lemma. By definition, the θ -span of a narrow ball is smaller than that of a wide ball. The lemma shows that this is also true, up to a constant factor, for the θ' -spans of the two balls. It actually establishes a slightly more general result.

Lemma 3.3 *Let $B_W, B_N \in \mathcal{B} \setminus \{B_0\}$, such that $\theta_{B_N} \leq \theta_{B_W}$. Then $\theta'_{B_N} = O(\theta'_{B_W}/\eta^2)$.*

Proof Let r_{B_W}, r_{B_N} be the respective radii of B_W, B_N , and let d_{B_W}, d_{B_N} be the respective distances of their centers, c_{B_W}, c_{B_N} , from c_0 . The fact that the angle between $c_0c_{B_W}$ and the z -axis is at least $\sin^{-1}\eta$, together with Claim 2 of Sect. 2 (applied to B_W), imply that

$$\theta_{B_W} = O(\min\{r_{B_W}/(\eta d_{B_W}), 1\}). \quad (2)$$

We also have

$$\theta_{B_W} \geq \theta_{B_N} = \Omega(\min\{r_{B_N}/d_{B_N}, 1\}). \quad (3)$$

Indeed, if the xy -projections of B_0, B_N intersect then $\theta_{B_N} = \Theta(1)$; otherwise we refer to Fig. 10 (with B_N instead of B_i) and observe that $\sin \theta_{B_N} = (r_0 + r_{B_N})/\hat{d}_{B_N} \geq r_{B_N}/d_{B_N}$.

¹²This is done in order to ensure that the surfaces of Γ^+, Γ^- be monotone.

¹³That is, balls of $\mathcal{B}_W^\Delta \cup \mathcal{B}_N^\Delta$ that “define” the WWN-lines appearing in the recursive sub-problem of τ .

¹⁴This explains why this new sub-recursion must be defined in the *new* coordinate frame. Working in the standard coordinate frame $(\theta, \kappa, \varphi)$ would not keep the surfaces of Γ^+, Γ^- $\theta\kappa$ -monotone.

The combination of (2) and (3) implies that

$$\min\{r_{B_N}/d_{B_N}, 1\} = O(\min\{r_{B_W}/(\eta d_{B_W}), 1\}). \quad (4)$$

Similarly to (2), the fact that the angle between $c_0 c_{B_N}$ and $\vec{v}$ is at least $\sin^{-1} \eta$ implies that $\theta'_{B_N} = O(\min\{r_{B_N}/(\eta d_{B_N}), 1\})$. Thus, (4) implies that $\theta'_{B_N} = O(\min\{r_{B_W}/(\eta^2 d_{B_W}), 1\})$.

Similarly to (3), we can show that $\theta'_{B_W} = \Omega(\min\{r_{B_W}/d_{B_W}, 1\})$, implying finally that $\theta'_{B_N} = O(\theta'_{B_W}/\eta^2)$. $\square$

For each ball $B \in \mathcal{B}_W^\Delta$ we construct the $\theta'\kappa'$ -monotone surface patches σ_B^+, σ_B^- , as defined above, and let Γ^+, Γ^- denote the corresponding collections of these patches. (Recall that these patches are defined for the fixed pair of arcs $\Delta' \subseteq \Delta$.)

We fix a sufficiently large constant parameter $r \geq 0$ and apply a secondary recursive decomposition of the problem, where the arc $\Delta' \subset \Delta$ is kept fixed and the decomposition is of the $\theta'\kappa'$ -domain. To do so, we extend the notions of wideness and narrowness to $\theta'\kappa'$ -regions (rather than θ -ranges), as follows. Let τ be any (connected) region within the $\theta'\kappa'$ -domain, and let θ'_τ denote the angular extent of the θ' -projection of τ . We call a ball $B \in \mathcal{B} \setminus \{B_0\}$ *wide* (resp., *narrow*) in τ if (i) B is stabbed by at least one line $\ell \in \mathcal{L}$ having $(\theta'(\ell), \kappa'(\ell)) \in \tau$, and (ii) B has θ' -span θ'_B larger than (resp., at most) θ'_τ . To distinguish between the two notions of wideness/narrowness, we refer to balls which are wide/narrow in the earlier sense as θ -wide or θ -narrow, and to balls which are wide/narrow in the new sense as θ' -wide or θ' -narrow, respectively.

At each recursive step we are given a (parent) pseudo-trapezoidal (and connected) region τ_0 in the $\theta'\kappa'$ -domain, a set $\mathcal{B}_N^{\tau_0}$ of n balls that are θ -narrow in Δ and θ' -narrow in τ_0 , and a set $\mathcal{B}_W^{\tau_0}$ of m balls that are θ -wide in Δ and contribute a surface patch to $\Gamma^+ \cup \Gamma^-$ which appears on the boundary of $\mathcal{F}(\mathcal{B})$ over τ_0 . The subsets of Γ^+, Γ^- consisting of these patches are denoted by $\Gamma_{\tau_0}^+, \Gamma_{\tau_0}^-$, respectively. Our goal is to bound the number of free lines $\ell \in \mathcal{L}$ such that $(\theta'(\ell), \kappa'(\ell)) \in \tau_0$, ℓ is tangent to two balls of $\mathcal{B}_W^{\tau_0}$ and to one ball in $\mathcal{B}_N^{\tau_0}$. Initially, τ_0 is the entire $\theta'\kappa'$ -domain, and $\mathcal{B}_N^{\tau_0} = \mathcal{B}_N^\Delta$; all these balls are indeed θ -narrow in Δ and θ' -narrow in τ_0 .

The parent region τ_0 is decomposed into $O(r^{2+\varepsilon})$ connected pseudo-trapezoidal cells τ , for any $\varepsilon > 0$, satisfying the following properties:

- (i) At most m/r surface patches of $\Gamma_{\tau_0}^+ \cup \Gamma_{\tau_0}^-$ appear on the boundary of $\mathcal{F}(\mathcal{B})$ over τ . We denote the set of all balls generating those patches by $\mathcal{B}_W^\tau$. Observe that any free WWN-line $\ell \in \mathcal{L}$, having $(\theta'(\ell), \kappa'(\ell)) \in \tau$, is tangent to two balls from $\mathcal{B}_W^\tau$.
- (ii) At most n/r balls from $\mathcal{B}_N^{\tau_0}$ remain θ' -narrow in τ . Let us denote the set of these balls as $\mathcal{B}_N^\tau$.

In more detail, our decomposition scheme is defined in two stages. First, we sample, randomly and uniformly, a collection of $br \log r$ surface patches from $\Gamma_{\tau_0}^+ \cup \Gamma_{\tau_0}^-$, for a sufficiently large constant $b > 0$, and treat each sampled patch as a partially defined function over the $\theta'\kappa'$ -domain. Denote by $\mathcal{M}^+$ (resp., $\mathcal{M}^-$) the portions within τ_0 of the minimization (resp., maximization) diagram of the surfaces sampled from Γ^+ (resp., Γ^-). Let $\mathcal{M}$ be the decomposition of τ_0 obtained by partitioning

Fig. 13 The pseudo-trapezoidal decomposition $\mathcal{M}'$ (at the first level of recursion, where τ_0 is the whole $\theta'\kappa'$ -domain). The edges of $\mathcal{M}$ are drawn *solid*, and the strip decomposition is drawn *dashed*

the overlay of $\mathcal{M}^+$ and $\mathcal{M}^-$ into pseudo-trapezoidal cells (of constant description complexity). The analysis in [4, 11] implies that $\mathcal{M}$ consists of $O(r^{2+\varepsilon})$ cells, for any $\varepsilon > 0$. Let τ be any fixed (pseudo-trapezoidal) cell in $\mathcal{M}$, and let Γ_τ^+ (resp., Γ_τ^-) denote the set of patches from $\Gamma_{\tau_0}^+$ (resp., $\Gamma_{\tau_0}^-$) that appear over τ on the boundary of the sandwich region below the lower envelope of $\Gamma_{\tau_0}^+$ and above the upper envelope of $\Gamma_{\tau_0}^-$.

The ε -net theory (see, e.g., [16]) guarantees that, for $b > 0$ sufficiently large, we have $|\Gamma_\tau^+| \leq m/(2r)$ and $|\Gamma_\tau^-| \leq m/(2r)$ (that is, τ satisfies condition (i)), with probability close to 1. This guarantees, with high probability, that $|\mathcal{B}_W^\tau| \leq m/r$, for any child subcell τ of τ_0 . We assume that our sample does indeed satisfy this property.

We further refine $\mathcal{M}$ as follows. We partition the θ' -projection of τ_0 into $O(r)$ circular arcs, each containing at most n/r endpoints of the arcs I_B^+ , I_B^- , for $B \in \mathcal{B}_N^\Delta$. We obtain a strip decomposition of the $\theta'\kappa'$ -domain by lifting each of the above circular arcs into a κ' -parallel strip. Let $\mathcal{M}'$ denote the overlay of $\mathcal{M}$ with the resulting strip decomposition. See Fig. 13. Clearly, each cell $\tau \in \mathcal{M}'$ satisfies $|\mathcal{B}_N^\tau| \leq n/r$, so the cells of $\mathcal{M}'$ satisfy both (i) and (ii).

The following lemma has been proved, in a different, albeit very similar context, in [5]. For the sake of completeness, we give a sketch of its proof for the setup considered here.

Lemma 3.4 *The decomposition $\mathcal{M}'$ consists of $O(r^{2+\varepsilon})$ pseudo-trapezoidal cells, for any $\varepsilon > 0$.*

Proof As the complexity of $\mathcal{M}$ is $O(r^{2+\varepsilon})$, for any $\varepsilon > 0$, it is sufficient to bound the number of vertices v in $\mathcal{M}'$ that are formed by the intersection between the edges of $\mathcal{M}$ and the κ' -parallel lines of the strip decomposition. Fix such a line $\ell^\perp$. Let $\tilde{\Gamma}^+$ (resp., $\tilde{\Gamma}^-$) denote the collection of $O(r \log r)$ arcs, each of which is formed by intersecting some patch of Γ^+ (resp., Γ^-) with the φ' -parallel (“vertical”) plane through $\ell^\perp$. Then each vertex v of $\mathcal{M}'$ of the above kind is the $\theta'\kappa'$ -projection of a vertex of the sandwich region which is enclosed between the lower envelope of $\tilde{\Gamma}^+$ and the upper envelope of $\tilde{\Gamma}^-$. The number of such vertices v is $O(\lambda_s(r) \log r)$, for an appropriate constant parameter s , for any fixed $\ell^\perp$ (see also [16]). The bound now follows by repeating this argument to each of the $O(r)$ lines $\ell^\perp$. $\square$

As the $\theta'\kappa'$ -domain is progressively cut into subcells τ , more and more balls of $\mathcal{B}_N^\Delta$ become θ' -wide in τ , while the size of $\mathcal{B}_N^\tau$ and $\mathcal{B}_W^\tau$ keeps decreasing. During

each step of the recursion we immediately dispose of any free WWN-lines tangent to the new θ' -wide balls from $\mathcal{B}_N^\Delta$ (that is, balls that were θ' -narrow in the parent cell τ_0 of τ). The number of these lines is bounded by the following lemma.

Lemma 3.5 *Let Δ be a circular arc on $\mathbb{S}^1$, and let τ be a (connected) region in the $\theta'\kappa'$ -domain. Let $\mathcal{B}_W^\tau \subset \mathcal{B} \setminus \{B_0\}$ be a collection of m balls which are θ -wide in Δ , and let $\mathcal{B}_N^\tau \subset \mathcal{B} \setminus \{B_0\}$ be another collection of n balls which are θ -narrow in Δ and θ' -wide in τ . Then the number of free lines $\ell \in \mathcal{L}$ which are charged to B_0 and satisfy: (i) $(\theta'(\ell), \kappa'(\ell)) \in \tau$, (ii) ℓ is tangent to two balls from $\mathcal{B}_W^\tau$ and to one ball from $\mathcal{B}_N^\tau$, and (iii) $\vec{v}$ is a good direction for ℓ , is $O((n+m)m^{1+\varepsilon})$, for any $\varepsilon > 0$.*

Informally, the lemma deals with isolated free lines $\ell \in \mathcal{L}$ tangent to two θ -wide balls (in Δ), and to a single θ -narrow ball which has now become θ' -wide (in τ). In addition, Lemma 3.3 implies that each of the two θ -wide balls is “sufficiently” θ' -wide in τ . Hence, we can bound the number of such WWN-lines by carefully applying Theorem 3.2 with respect to the new coordinate frame $(\theta', \kappa', \varphi')$.

Proof We partition the set of all balls $B \in \mathcal{B}_N^\tau$ into $O(\lceil n/m \rceil)$ sets of size $O(m)$ each. For each resulting set $\mathcal{B}'$ we establish the bound $O(m^{2+\varepsilon})$, for any $\varepsilon > 0$, on the number of free WWN-lines $\ell \in \mathcal{L}$ as in the lemma, which are tangent to two balls from $\mathcal{B}_W^\Delta$, and to one from $\mathcal{B}'$. Summing over all subsets $\mathcal{B}'$, and observing that all lines under consideration are counted in this sum, the bound asserted in the lemma follows.

Fix a subset $\mathcal{B}'$. Each ball $\hat{B} \in \mathcal{B}'$ satisfies $\theta'_B \geq \theta'_\tau$. On the other hand, each ball $B \in \mathcal{B}_W^\tau \subset \mathcal{B}_W^\Delta$ satisfies $\theta_B \geq \theta_\Delta \geq \theta_{\hat{B}}$, and thus, by Lemma 3.3, we have $\theta'_B = \Omega(\eta^2 \theta'_\tau)$. Hence $\theta'_B = \Omega(\eta^2 \theta'_\tau)$, for every $B \in \mathcal{B}_W^\tau$.

We partition τ , by κ' -parallel lines, into c/η^2 pseudo-trapezoidal subcells $\mu \subset \tau$, where c is a sufficiently large constant, such that $\theta'_\mu = O(\eta^2 \theta'_\tau)$, for every μ . Then, with an appropriate choice of c , for each $B \in \mathcal{B}_W^\tau \cup \mathcal{B}'$, either B is θ' -wide in μ , or $I_B^{'+} \cup I_B^{'-}$ is disjoint from the θ' -projection Δ'_μ of μ . Balls of the second kind can be ignored because they do not induce any free lines with $(\theta'(\ell), \kappa'(\ell)) \in \mu$. We now apply Theorem 3.2 to the set $\mathcal{B}_W^\tau \cup \mathcal{B}'$ (and the arc $\Delta'_\mu \subset \mathbb{S}^1$). In this application we regard $\vec{v}$ as the positive z -direction. Then the balls of $\mathcal{B}_W^\tau \cup \mathcal{B}'$ are wide, in the original sense, with respect to Δ'_μ . In this new setting, Theorem 3.2 implies that the number of free lines in $\mathcal{L}$, charged to B_0 and admitted by the balls of $\mathcal{B}' \cup \mathcal{B}_W^\tau$, such that $(\theta'(\ell), \kappa'(\ell)) \in \mu$ (and hence $\theta'(\ell) \in \Delta'_\mu$), and $\vec{v}$ is good for ℓ , is $O(m^{2+\varepsilon})$, for any $\varepsilon > 0$.

Recall that in order to apply Theorem 3.2 in the new context, we must assume that, for every $B \in \mathcal{B}_W^\tau \cup \mathcal{B}'$, the angle between $c_0 c_B$ and the $\vec{v}$ -direction is at least $\sin^{-1} \eta$. This assumption can be made because of the additional pruning step imposed on the set of good directions in Sect. 2 (see Remark 2 at the end of that section), and because $\vec{v}$ is assumed to be a good direction, belonging to the set constructed in the pruning procedure of Sect. 2, for all isolated free lines $\ell \in \mathcal{L}$ counted by the current lemma. This allows us to consider in $\mathcal{B}_W^\tau \cup \mathcal{B}'$ only the balls $B \in \mathcal{B} \setminus \{B_0\}$ for which the angle between $c_0 c_B$ and the $\vec{v}$ -direction is at least $\sin^{-1} \eta$. $\square$

We now return to the “global” sub-problem involving a fixed arc Δ (and a sub-arc $\Delta' \subseteq \mathbb{S}^1$). Let $F_1(n, m)$ denote the maximum number of free WNN-lines $\ell \in \mathcal{L}$ charged to B_0 and tangent to one ball from $\mathcal{B}_N^\Delta$ and two balls from $\mathcal{B}_W^\Delta$, such that $\vec{v}$ is a good direction for ℓ , over all possible collections $\mathcal{B}_W^\Delta \subset \mathcal{B} \setminus \{B_0\}$ and $\mathcal{B}_N^\Delta \subset \mathcal{B} \setminus \{B_0\}$ of m (resp., n) θ -wide (resp., θ -narrow) balls in Δ . Then $F_1(n, m)$ satisfies the following recurrence.

$$F_1(n, m) = \begin{cases} O(r^{2+\varepsilon_0})F_1(\frac{n}{r}, \frac{m}{r}) + O((n+m)m^{1+\varepsilon_0}), & \min\{n, m\} \geq r, \\ O(n+m^2), & \min\{n, m\} < r, \end{cases} \quad (5)$$

where $r > 0$ is a sufficiently large constant, $\varepsilon_0 > 0$ is an arbitrarily small constant, and the constants of proportionality depend on ε_0 , and, in the two non-recursive terms, also on r . (The second bound is the trivial bound $O(nm^2)$, specialized for the case $\min\{m, n\} < r$) The recurrence solves to $F_1(n, m) = O((n+m)m^{1+\varepsilon})$, for any $\varepsilon > \varepsilon_0$. Hence, the bound holds for any $\varepsilon > 0$, by choosing a smaller value for ε_0 . Thus, summing over all good directions and subarcs Δ' , we obtain the following theorem.

Theorem 3.6 *The number of free WNN-lines of $\mathcal{L}$ (which are charged to B_0) in Δ , with respect to a set $\mathcal{B}_W^\Delta$ of m θ -wide balls and a set $\mathcal{B}_N^\Delta$ of n θ -narrow balls, is $O((n+m)m^{1+\varepsilon})$, for any $\varepsilon > 0$.*

3.5 The Number of Free WNN-Lines

We follow a similar approach to the one in the case of free WNN-lines. We have fixed arcs $\Delta' \subseteq \Delta$. We apply the same recursive decomposition in the $\theta'\kappa'$ -domain, and extend the notions of wideness and narrowness to the θ' -coordinate, as above. At each recursive step we are given a (parent) pseudo-trapezoidal (and connected) region τ_0 of the $\theta'\kappa'$ -domain, a set $\mathcal{B}_N^{\tau_0}$ of n balls that are θ -narrow in Δ and θ' -narrow in τ_0 , and a set $\mathcal{B}_W^{\tau_0}$ of m balls that are θ -wide in Δ . Our goal is to bound the number of free lines $\ell \in \mathcal{L}$ charged to B_0 , such that $(\theta'(\ell), \kappa'(\ell)) \in \tau_0$, ℓ is tangent to two balls of $\mathcal{B}_N^{\tau_0}$ and to one ball in $\mathcal{B}_W^{\tau_0}$ (and $\vec{v}$ is a good direction for ℓ). We denote by $\Gamma_{\tau_0}^+$, $\Gamma_{\tau_0}^-$ the subsets of Γ^+ , Γ^- , respectively, of those patches that appear on the boundary of $\mathcal{F}(\mathcal{B})$ over τ_0 . We construct the decomposition $\mathcal{M}'$ of τ_0 as in the case of WNN-lines. We recall its properties. It consists of $O(r^{2+\varepsilon})$ connected pseudo-trapezoidal cells, for any $\varepsilon > 0$, and satisfies the following, for each cell τ .

- (i) At most m/r surface patches of $\Gamma_{\tau_0}^+ \cup \Gamma_{\tau_0}^-$ appear on the boundary of $\mathcal{F}(\mathcal{B})$ over τ . We denote the set of all balls generating those patches as $\mathcal{B}_W^\tau$. Observe that any free WNN-line $\ell \in \mathcal{L}$, having $(\theta'(\ell), \kappa'(\ell)) \in \tau$, is tangent to one ball from $\mathcal{B}_W^\tau$.
- (ii) At most n/r balls from $\mathcal{B}_N^{\tau_0}$ remain θ' -narrow in τ . Let us denote the set of these balls as $\mathcal{B}_N^\tau$.

As the $\theta'\kappa'$ -domain is progressively cut into subcells τ , more and more balls of $\mathcal{B}_N^\Delta$ become θ' -wide in the new subcells τ , while the sizes of the corresponding sets $\mathcal{B}_N^\tau$ and $\mathcal{B}_W^\tau$ keep decreasing (by properties (i) and (ii) above). During each step of the recursion we immediately dispose of any free WNN-lines tangent to the new

θ' -wide balls from $\mathcal{B}_N^\Delta$ (that is, balls that were θ' -narrow in the parent cell τ_0 of τ). The number of these lines is bounded by the following lemma.

Lemma 3.7 *Let Δ be a circular arc on $\mathbb{S}^1$, and let τ be a (connected) region in the $\theta'\kappa'$ -domain. Let $\mathcal{B}_W^\tau \subset \mathcal{B} \setminus \{B_0\}$ be a collection of m balls which are θ -wide in Δ , and let $\mathcal{B}_N^\tau \subset \mathcal{B} \setminus \{B_0\}$ be another collection of n balls which are θ -narrow in Δ . Then the number of free lines $\ell \in \mathcal{L}$, which are charged to B_0 and satisfy: (i) $(\theta'(\ell), \kappa'(\ell)) \in \tau$, (ii) ℓ is tangent to one ball from $\mathcal{B}_W^\tau$ and to two balls from $\mathcal{B}_N^\tau$, at least one of which is θ' -wide in τ , and (iii) $\vec{v}$ is a good direction for ℓ , is $O((n+m)n^{1+\varepsilon})$, for any $\varepsilon > 0$.*

Proof We partition $\mathcal{B}_W^\tau$ into $O(\lceil m/n \rceil)$ sets of size $O(n)$ each. For each resulting set $\mathcal{B}'$ we establish the bound $O(n^{2+\varepsilon})$, for any $\varepsilon > 0$, on the number of free WNN-lines $\ell \in \mathcal{L}$ as in the lemma, which are tangent to one ball from $\mathcal{B}'$, and to two balls from $\mathcal{B}_N^\tau$, at least one of which is θ' -wide in τ . Summing over all subsets $\mathcal{B}' \subset \mathcal{B}_W^\tau$, and observing that all WNN-lines under consideration are counted in this sum, the bound asserted in the lemma follows.

Fix a subset $\mathcal{B}'$, as above. As before, we have $\theta'_B = \Omega(\eta^2\theta'_\tau)$, for every $B \in \mathcal{B}'$. We partition τ , by κ' -parallel lines, into c/η^2 pseudo-trapezoidal subcells $\mu \subset \tau$, where c is a sufficiently large constant, such that $\theta'_\mu = O(\eta^2\theta'_\tau)$. Fix such a subcell μ . Then, with an appropriate choice of c , for each $B \in \mathcal{B}'$, either B is θ' -wide in μ , or $I_B^{'+} \cup I_B^{'-}$ is disjoint from the θ' -projection Δ'_μ of μ . Let $\mathcal{B}''$ denote the subset of the balls of $\mathcal{B}'$ of the former kind, and let $\mathcal{B}'''$ denote the subset of those balls in $\mathcal{B}_N^\tau$ which are θ' -wide in μ . Then any free WNN-line ℓ that we consider here, for which $(\theta'(\ell), \kappa'(\ell)) \in \mu$, is tangent to one ball of $\mathcal{B}'$, one ball of $\mathcal{B}'''$, and one ball of $\mathcal{B}_N^\tau$. By construction, the first two balls are θ' -wide in μ , and the third ball is either θ' -wide or θ' -narrow. We can therefore apply either Theorem 3.2 (to the arc $\Delta'_\mu \subset \mathbb{S}^1$ and the set $\mathcal{B}'' \cup \mathcal{B}'''$) or Theorem 3.6 (to the arc $\Delta'_\mu \subset \mathbb{S}^1$ and the sets $\mathcal{B}'' \cup \mathcal{B}'''$, $\mathcal{B}_N^\tau \setminus \mathcal{B}'''$), regarding the $\theta'\kappa'$ -frame as the standard frame (see a discussion of this step in the proof of Lemma 3.5), to conclude that the number of free WNN-lines under consideration is $O(n^{2+\varepsilon})$, for any $\varepsilon > 0$. This completes the proof of Lemma 3.7. $\square$

We now return to the “global” sub-problem involving a fixed arc Δ (and a sub-arc $\Delta' \subseteq \mathbb{S}^1$). Let $F_2(n, m)$ denote the maximum number of free WNN-lines $\ell \in \mathcal{L}$ charged to B_0 , and tangent to one ball from $\mathcal{B}_W^\Delta$ and two balls from $\mathcal{B}_N^\Delta$, such that $\vec{v}$ is a good direction for ℓ , over all possible collections $\mathcal{B}_W^\Delta$ and $\mathcal{B}_N^\Delta$ of m (resp., n) θ -wide (resp., θ -narrow) balls in Δ . Then $F_2(n, m)$ satisfies the following recurrence.

$$F_2(n, m) = \begin{cases} O(r^{2+\varepsilon_0})F_2(\frac{n}{r}, \frac{m}{r}) + O((n+m)n^{1+\varepsilon_0}), & \min\{n, m\} \geq r, \\ O(n^2 + m), & \min\{n, m\} < r, \end{cases} \quad (6)$$

where $r > 0$ is a sufficiently large constant, $\varepsilon_0 > 0$ is an arbitrarily small constant, and the constants of proportionality depend on ε_0 , and, in the two non-recursive terms, also on r . (The second bound is the trivial bound $O(n^2m)$, specialized to the case $\min\{n, m\} < r$.) The recurrence solves to $F_2(n, m) = O((n+m)n^{1+\varepsilon})$, for any $\varepsilon > \varepsilon_0$, and so it holds for any $\varepsilon > 0$, by choosing $\varepsilon_0 < \varepsilon$. Thus, summing over all directions $\vec{v}$ and subarcs Δ' , we obtain the following theorem.

Theorem 3.8 *The number of free WNN-lines of $\mathcal{L}$ (which are charged to B_0) in Δ , with respect to a set $\mathcal{B}_W^\Delta$ of m θ -wide balls and a set $\mathcal{B}_N^\Delta$ of n θ -narrow balls, is $O((n+m)n^{1+\varepsilon})$, for any $\varepsilon > 0$.*

Returning to the overall recursion scheme, we apply Theorems 3.2, 3.6, and 3.8 to bound the overall number of new free WWW-, WWN-, and WNN-lines that are disposed of while processing a single θ -arc $\Delta \subset \mathbb{S}^1$. Let n denote the number of balls $B \in \mathcal{B} \setminus \{B_0\}$ which are θ -narrow at the parent arc Δ_0 of Δ . Then $|\mathcal{B}_W^\Delta| + |\mathcal{B}_N^\Delta| \leq n$, and the preceding analysis implies that the number of lines disposed of at Δ is $O(n^{2+\varepsilon})$, for any $\varepsilon > 0$. Hence, the number of isolated free tangent lines of $\mathcal{L}$ admitted by $\mathcal{B}$, which are charged to B_0 , is bounded by the function $F(n)$ satisfying the recurrence (1). As we have shown, $F(n) = O(n^{2+\varepsilon})$, for any $\varepsilon > 0$. This completes the proof of Theorem 2.3, and thus also of the main Theorem 2.4. $\square$

4 Discussion

4.1 Finding Free Lines

Unlike the analysis of [1], our proof of Theorem 2.4 is constructive, and leads to a deterministic algorithm which finds all isolated free lines in $O(n^{3+\varepsilon})$ time, for any $\varepsilon > 0$. For each ball $B_0 \in \mathcal{B}$, we consider the space $\mathcal{L} = \mathcal{L}(B_0)$ of lines tangent to B_0 , and we compute in $O(n^{2+\varepsilon})$ time, for any $\varepsilon > 0$, a superset $A_{\mathcal{L}}$ of isolated tangent lines in $\mathcal{L}$ which contains all isolated free lines charged to B_0 . We can use the data structure of [13] which supports ray shooting amid the spheres of $\mathcal{B}$ in polylogarithmic time, and requires $O(n^{3+\varepsilon})$ storage and preprocessing time, for any $\varepsilon > 0$, to find all free lines in $A_{\mathcal{L}}$.

Computing $A_{\mathcal{L}}$ We closely follow the overall analysis (and the notation) of Sect. 3 which recurses on the circular arc $\Delta_0 \subset \mathbb{S}^1$ and the set $\mathcal{B}_N^{\Delta_0} \subseteq \mathcal{B} \setminus \{B_0\}$ of n balls which are θ -narrow in Δ_0 (starting with $\Delta_0 = \mathbb{S}^1$ and $\mathcal{B}_N^{\Delta_0} = \mathcal{B} \setminus \{B_0\}$). At each (non-terminal) stage of the recursion we decompose Δ_0 into $2r$ subarcs Δ , in $O(n)$ time. For each constructed subarc $\Delta \subseteq \Delta_0$ we determine the collections $\mathcal{B}_W^\Delta, \mathcal{B}_N^\Delta$, as defined above, in a total of $O(nr)$ time, and find three supersets of isolated tangent lines in $\mathcal{L}$ containing respectively all free WWW-, WWN-, and WNN-lines in Δ (that are charged to B_0). Below we describe how to perform the latter tasks in $O(n^{2+\varepsilon})$ time, for any $\varepsilon > 0$, implying that the overall running time satisfies inequality (1), and is therefore $O(n^{2+\varepsilon})$, for any $\varepsilon > 0$. (At the bottom of the recursion we apply a brute-force procedure to compute all the relevant free tangent lines.)

Finding WWW-Lines Following the analysis of Sect. 3, we consider an appropriate collection of $O(1)$ reference directions $\vec{v} \in \mathbb{S}^2$. For each subsequent choice of $\vec{v}$ we find a superset containing all free WWW-lines in Δ having $\vec{v}$ as a good direction (which are charged to B_0). We partition Δ into $O(1/\delta)$ subarcs $\Delta' \subseteq \Delta$, as in Sect. 3. For each resulting subarc Δ' we work in the $(\theta', \kappa', \varphi')$ -frame induced by $\vec{v}$ and construct the corresponding collections Γ^+, Γ^- of $\theta'\kappa'$ -monotone surfaces, in $O(n)$

time. Then we apply the algorithm of [4] to compute, in $O(n^{2+\varepsilon})$ time, for any $\varepsilon > 0$, all lines corresponding to vertices of the sandwich region which lies below the lower envelope of Γ^+ and above the upper envelope of Γ^- . Theorem 3.2 ensures that in this way we find all free WWW-lines in Δ' that are charged to B_0 and have $\vec{v}$ as a good direction.

Finding WWN-Lines Without loss of generality, we fix the quantities $\Delta' \subseteq \Delta \subseteq \Delta_0$, $\vec{v} \in \mathbb{S}^2$, Γ^+ and Γ^- , as above. Closely following the analysis (and the notation) for the number of free WWN-lines in Δ' that are charged to B_0 and have $\vec{v}$ as a good direction, we recurse on the $\theta'\kappa'$ -region τ_0 , the collection $\mathcal{B}_N^{\tau_0} \subseteq \mathcal{B}_N^\Delta$ of n balls which are θ' -narrow in τ_0 (and θ -narrow in Δ), and the collection $\mathcal{B}_W^{\tau_0}$ of m balls which are θ -wide in Δ .

We fix r to be a sufficiently large constant. At each (non-terminal) stage of the recursion we select $br \log r$ surface patches of $\Gamma_{\tau_0}^+ \cup \Gamma_{\tau_0}^-$ as in Sect. 3, and construct, in naive $O(r^4 m) = O(m)$ time, the pseudo-trapezoidal decomposition $\mathcal{M}$ of τ_0 , each of whose cells τ is such that at most m/r surface patches of $\Gamma_{\tau_0}^+ \cup \Gamma_{\tau_0}^-$ appear on the boundary of $\mathcal{F}(\mathcal{B})$ over τ . Here, instead of sampling $br \log r$ patches at random, we apply the deterministic method of [12] for finding $(1/r)$ -nets in general range spaces. This can be done in $O(mr^D) = O(m)$ time, for an appropriate constant $D > 0$. To construct the refined decomposition $\mathcal{M}'$ as in Sect. 3, we overlay $\mathcal{M}$ with an arrangement of $2r\kappa'$ -parallel lines, in naive $O(n) + O(r|\mathcal{M}|)$ time. Finally, for each of the $O(r^{2+\varepsilon})$ cells $\tau \in \mathcal{M}'$, we construct the respective collections $\mathcal{B}_W^\tau$ and $\mathcal{B}_N^\tau$, in a total of $O((n+m)r^{2+\varepsilon}) = O(n+m)$ time.

For each pseudo-trapezoidal region $\tau \in \mathcal{M}'$ we construct a superset of isolated tangent lines containing all free WWN-lines that are tangent to at least one ball of $\mathcal{B}_N^{\tau_0}$ that is θ' -wide in τ . The reduction in Lemma 3.5, based on partitioning $\mathcal{B}_N^{\tau_0}$ into subsets, allows us to apply the previously described algorithm for finding a superset containing free WWN-lines. Thus, the overall time required to compute a superset containing all free WWN-lines in Δ satisfies the inequality in (5), and thus has the same asymptotic bound on its running time.

Finding WNN-Lines We use the same algorithm as the one used above for WWN-lines, with the following modification. For each pseudo-trapezoidal region $\tau \in \mathcal{M}'$ we construct a superset of isolated tangent lines containing all free WNN-lines that are tangent to at least one ball of $\mathcal{B}_N^{\tau_0}$ that is θ' -wide in τ . The reduction in Lemma 3.7, based on partitioning $\mathcal{B}_N^{\tau_0}$ into subsets, allows us to use the previously described algorithms for constructing supersets containing free WWN-lines and free WWW-lines. Thus, the overall time required to compute a superset containing all free WWN-lines in Δ satisfies the inequality in (6), and thus has the same asymptotic bound on its running time.

We thus have the following result.

Theorem 4.1 *Let $\mathcal{B}$ be a collection of n balls in $\mathbb{R}^3$. (i) The collection of all isolated free lines in $\mathcal{F}(\mathcal{B})$ can be constructed in $O(n^{3+\varepsilon})$ time, for any $\varepsilon > 0$. (ii) A discrete representation of the boundary of the space $\mathcal{F}(\mathcal{B})$ of all free lines, i.e., the set of all lines that miss $\mathcal{B}$, can be constructed in $O(n^{3+\varepsilon})$ time, for any $\varepsilon > 0$.*

Part (ii) of the theorem follows from part (i) by a fairly standard procedure, which connects the isolated free lines, which are vertices of $\mathcal{F}(\mathcal{B})$, by edges,¹⁵ then connects the edges along the boundaries of 2-faces, and then glues together the 2-faces along the boundaries of the 3-cells of $\partial\mathcal{F}(\mathcal{B})$. The actual details of the construction are somewhat subtle but fairly routine, and we omit them in this paper. See [16, Sects. 7.4.2 and 7.4.3] and [2] for constructions of a similar nature.

4.2 Union of Geometric Objects and Free Lines

Our analysis reduces the problem of bounding the number of isolated free tangent lines in a collection of balls to the problem of bounding the number of certain vertices on the boundary of the union of regions of a certain kind in parametric 3-space, showing that this latter quantity is only nearly quadratic.¹⁶ Hence, a side interpretation of our result is that it provides yet another class of objects in 3-space whose union has nearly quadratic complexity (admittedly, our objects are rather “unnatural”), adding to the growing body of such instances studied in the past decade (see [3] for a recent comprehensive survey.) With this interpretation, it is not surprising that our analysis parallels that of Ezra [5], which also bounds the complexity of the union of objects (arbitrary cylinders) in 3-space. It would be interesting to find other applications of this general machinery to new instances of union of objects in 3-space.

4.3 Overlays of Minimization Diagrams in $\mathbb{R}^2$ and in $\mathbb{R}^3$

The analysis of [1], for the case of unit balls, uses an upper bound of $O(n^{3+\varepsilon})$, for any $\varepsilon > 0$, on the complexity of the overlay of two minimization diagrams in $\mathbb{R}^3$, whose proof in [11] requires a fairly intricate analysis. In contrast, our (more general) result only needs the simpler bound of $O(n^{2+\varepsilon})$, for any $\varepsilon > 0$, on the complexity of the overlay of two minimization diagrams in $\mathbb{R}^2$, given in [4].

4.4 Open Problems

The construction of [7] suggests that our bound on the number of isolated free tangents to a collection of n arbitrary balls is nearly tight in the worst case. However, this is still open for *unit* balls, because the construction of [7] uses balls of different radii, whereas the best known lower bound for unit balls, as given in [1], is only $\Omega(n^2)$. Hence, the problem of finding a sharp bound on the maximum number of isolated free tangents to a collection of n unit balls in $\mathbb{R}^3$ is still open.

Acknowledgements Work on this paper has been partially supported by Grants 155/05, 975/06 and 338/09 from the Israel Science Fund. This is part of the author’s Ph.D. Dissertation, prepared under the supervision of Haim Kaplan and Micha Sharir.

¹⁵Recall that these edges are contained in $O(n^3)$ curves, each representing lines tangent to some three balls of $\mathcal{B}$.

¹⁶To be pedantic, we only count vertices (lines) charged to B_0 . We therefore need a more careful definition of the regions whose union is being considered.

I would like to thank my advisor Micha Sharir for helpful discussions concerning the results of the paper, and, in particular, for his invaluable help in the proof of Theorem 2.2 and in the preparation of the paper.

In addition, I would like to thank the anonymous referee for valuable suggestions that greatly helped to improve the presentation.

References

1. Agarwal, P.K., Aronov, B., Koltun, V., Sharir, M.: On lines avoiding unit balls in three dimensions. *Discrete Comput. Geom.* **34**, 231–250 (2005)
2. Agarwal, P.K., Aronov, B., Sharir, M.: Computing envelopes in four dimensions with applications. *SIAM J. Comput.* **26**(6), 1714–1732 (1997)
3. Agarwal, P.K., Pach, J., Sharir, M.: State of the union—of geometric objects. In: Goodman, J.E. et al. (eds.) *Surveys in Discrete and Computational Geometry: Twenty Years Later*. Contemporary Mathematics, vol. 453, pp. 9–48. Am. Math. Soc., Providence (2008)
4. Agarwal, P.K., Schwarzkopf, O., Sharir, M.: The overlay of lower envelopes in 3-space and its applications. *Discrete Comput. Geom.* **15**, 1–13 (1996)
5. Ezra, E.: On the union of cylinders in three dimensions. *Discrete Comput. Geom.* **45**(1), 45–64 (2011)
6. Ezra, E., Sharir, M.: Counting and representing intersections among triangles in three dimensions. *Comput. Geom., Theory Appl.* **32**, 196–215 (2005)
7. Glisse, M., Lazard, S.: On the complexity of sets of free lines and line segments among balls in three dimensions. In: *Proc. 26th Annu. Symp. Comput. Geom.*, pp. 48–57 (2010)
8. Hodge, W., Pedoe, D.: *Methods in Algebraic Geometry*. Cambridge University Press, Cambridge (1952)
9. Kaplan, H., Rubin, N., Sharir, M.: Linear data structures for fast ray-shooting amidst convex polyhedra. *Algorithmica* **55**(2), 283–310 (2009)
10. Kaplan, H., Rubin, N., Sharir, M.: Line transversals of convex polyhedra in $\mathbb{R}^3$. *SIAM J. Comput.* **39**(7), 3283–3310 (2010)
11. Koltun, V., Sharir, M.: The partition technique for overlays of envelopes. *SIAM J. Comput.* **32**(4), 841–863 (2003)
12. Matoušek, J.: Approximations and optimal geometric divide-and-conquer. *J. Comput. Syst. Sci.* **50**(2), 203–208 (1995)
13. Mohabab, S., Sharir, M.: Ray shooting amidst spheres in 3 dimensions and related problems. *SIAM J. Comput.* **26**, 654–674 (1997)
14. Pach, J., Sharir, M.: *Combinatorial Geometry and Its Algorithmic Applications: The Alcalá Lectures*. Am. Math. Soc., Providence (2009)
15. Pellegrini, M.: On lines missing polyhedral sets in 3-space. *Discrete Comput. Geom.* **12**, 203–221 (1994)
16. Sharir, M., Agarwal, P.K.: *Davenport–Schinzel Sequences and Their Geometric Applications*. Cambridge University Press, Cambridge (1995)
17. Wenger, R.: Helly-type theorems and geometric transversals. In: Goodman, J.E., O’Rourke, J. (eds.) *Handbook of Discrete and Computational Geometry*, 2nd edn., pp. 73–96. Chapman & Hall/CRC Press, London/Boca Raton (2004)