On the size of the sphere of influence graph

Michael A. Soss

School of Computer Science

McGill University

Montréal, Canada

July 1998

A thesis submitted to the
Faculty of Graduate Studies and Research
in partial fulfillment of the requirements for
the degree of Master of Science

0-612-44285-3

Abstract

Let V be a set of distinct points in some metric space. We draw the following proximity graph on V: For each point $x \in V$, let s_x be the open ball centered at x with radius from x to its nearest neighbour. Then (a, b) is an edge if and only if s_a and s_b intersect. This graph is known as the sphere of influence graph of the point set V. In this thesis, we demonstrate that in the d-dimensional infinite-order Minkowski metric, no sphere of influence graph of n vertices contains $(2^{2d-1} - 2^{d-1})n$ edges or more. We also prove an asymptotic lower bound of $(2^{d+2} - 3d - 4)n/9$ on the maximum size of the graph. Lastly, we demonstrate an upper bound of $15n$ on the size of the sphere of influence graph of n vertices in the Euclidean plane.

Résumé

Soit V un ensemble de points dans un espace métrique. Nous créons le graphe de proximité suivant sur V: pour chaque point $x \in V$, soit s_x la boule ouverte de centre x et de rayon égale à la distance entre x et son voisin le plus proche. Nous créons aussi l'arc (a, b) du graphe si et seulement si s_a et s_b s'intersectent. Ce graphe s'appelle le graphe d'influence de sphères de l'ensemble V. Dans ce mémoire, nous prouvons que dans l'espace de dimension d selon la métrique de Minkowski, aucun graphe d'influence de sphères ne peut avoir $(2^{2d-1} - 2^{d-1})n$ arcs ou plus. Nous démontrons aussi que la borne inférieure asymptotique de la taille maximale est au moins de $(2^{d+2} - 3d - 4)n/9$. Finalement, nous prouvons que dans le plan euclidien, aucun graphe d'influence de sphères de n nœuds ne peut contenir plus de $15n$ arcs.

Statement of Originality

Except for the introduction, the appendix, and results whose authors have been cited where first introduced, all elements of this thesis are original contributions to knowledge. Assistance has been received only where mentioned in the acknowledgements.

Acknowledgements

I would first like to thank Godfried Toussaint for supervising this research and for introducing me to the field of computational geometry.

Next I wish to thank Jit Bose and several friends in the Computational Geometry group at McGill for the many discussions during my research: Therese Biedl, Sylvain Lazard, Tallman Nkgau, Sue Whitesides, and especially Steve Robbins, whose conversations proved most helpful whenever my proofs were not turning out exactly as I desired.

No thesis could ever be completed without the help of the departmental secretaries, Franca, Josie, Judy, Lise, Lucy, Vera, and Vicki, and I offer my appreciation.

Lastly, I thank my parents for their support and for the many things that have nothing to do with sphere of influence graphs.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

Suppose you gave a six-year-old child Figure 1.1a on a sheet of paper and asked her to draw what she thought it represented. How would she decide on the underlying structure of these points? She might connect a few points here and there and end up with something similar to Figure 1.1b.

If so, on the sheet of paper now is a set of points and a set of lines connecting those points. She has unknowingly drawn a *graph*. In formal terms, a graph is set of points, called *vertices*, and a set of *edges*, each connecting a pair of vertices. For example, in Figure 1.2, there are edges between p and q and between q and r. We

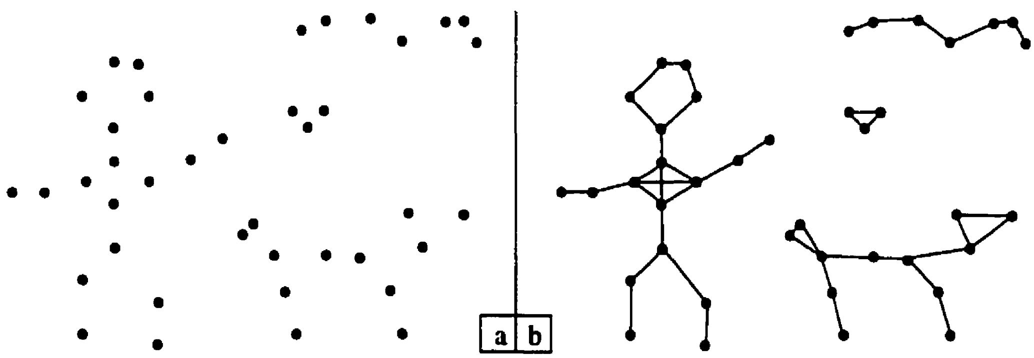

Figure 1.1: (a) A set of points and (b) a possible interpretation.

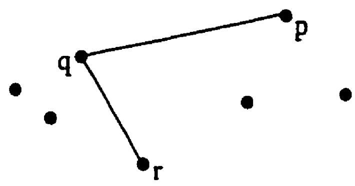

Figure 1.2: A graph.

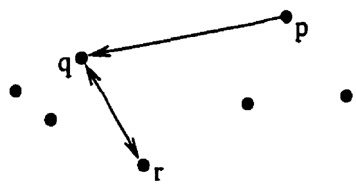

Figure 1.3: A directed graph.

then say that p and q are *neighbours*, as are q and r, since they share an edge.

We will also make reference to *directed edges*, which are ordered pairs of vertices. In Figure 1.3, there are edges from p to q, from q to r, and from r to q. A graph consisting of a set of vertices and a set of directed edges is known as a *directed graph.*

Let's look again at the child's artwork in Figure 1.1b. In drawing the edges, she didn't just arbitrarily choose pairs of vertices and connect them. Instead she connected those vertices which she thought were close to each other. This type of graph, in which vertices are connected by an edge if they satisfy some condition of closeness, is called a *proximity graph.*

Since we can define closeness in many different ways, there are several types of proximity graphs. We will not be describing in great detail the various graphs; however, the survey paper by Jaromczyk and Toussaint [JT92] contains an excellent discussion. The most famous proximity graph, the *minimal spanning tree* (MST), is demonstrated in Figure 1.4a. The MST is the connected graph which uses the minimum total edge length (which always results in one fewer edge than the number of vertices). By *connected graph,* we mean that there is path of edges between any two vertices. In the *relative neighbourhood graph* (RNG), shown in Figure 1.4b, two points are connected with an edge if their *lune* contains no other vertices of the graph. The lune of two points x and y, shown in Figure 1.5a, is defined as the intersection of two spheres of radius $dist(x, y)$, one centered at x and the other at y. In the *Gabriel graph* (GG), shown in Figure 1.4c, two points x and y are connected with an edge if their *diametral sphere*, the sphere of diameter $\overline{xy}$, contains no other vertices of the graph. Figure 1.4d contains another proximity graph, the *Delaunay triangulation* (DT). This graph connects two points x and y if there exists some sphere which has the chord $\overline{xy}$ and which contains no other vertices of the graph. All four graphs are *planar*, meaning no two edges cross, and connected. Furthermore, they share an interesting relationship; for any point set, MST $\subseteq$ RNG $\subseteq$ GG $\subseteq$ DT. This hierarchy makes these four graphs a nice mathematical tool for detecting the underlying structures of dot patterns. If one graph is too sparse for a particular purpose, the next graph in the relationship may prove more useful.

In this thesis, we will focus on a different kind of proximity graph, using the following rule to decide closeness. For each vertex v, we draw a circle centered at v with a radius equal to the distance to its nearest neighbour. This is shown in Figure 1.6a. For each pair of circles that properly intersect, we connect the two corresponding vertices with an edge. We say that two circles *properly intersect* if their intersection has a nonzero area (or volume, in higher dimensions), meaning that they intersect at more than just their boundaries. We call the circles of Figure 1.6a the *spheres of influence* of the vertices, and the resulting graph in Figure 1.6b is known as the *sphere of influence graph* (SIG). Since the SIG is necessarily neither planar nor connected, it does not fit into the MST $\subseteq$ RNG $\subseteq$ GG $\subseteq$ DT hierarchy mentioned above.

In the 1980's, Godfried Toussaint proposed the sphere of influence graph as a geometric tool for capturing the underlying structures of dot patterns [Tou80, Tou81,

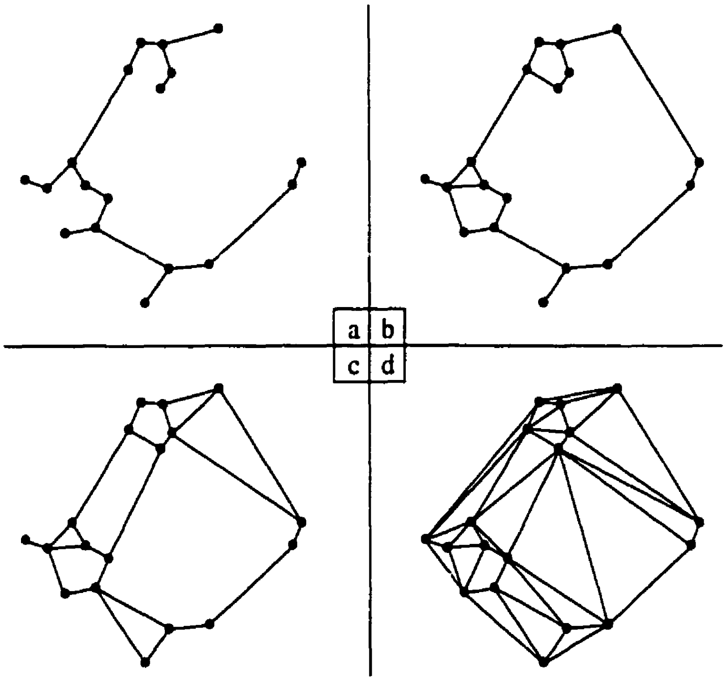

Figure 1.4: (a) A minimal spanning tree, (b) a relative neighbourhood graph, (c) a Gabriel graph, and (d) a Delaunay triangulation.

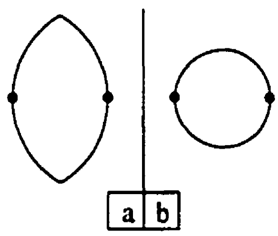

Figure 1.5: (a) A lune determined by two points and (b) their diametral sphere.

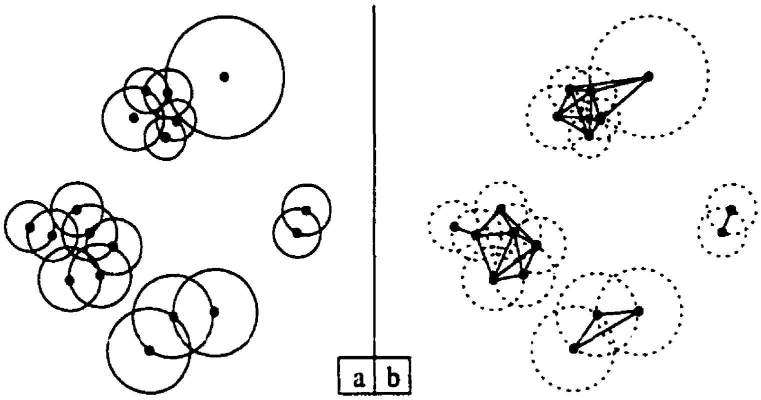

Figure 1.6: A set of points, (a) its spheres of influence, and (b) its sphere of influence graph.

Tou88]. We will concern ourselves with determining the maximum size of this graph. In other words, given a set of n vertices, what is the maximum number of edges the resulting SIG can have?

As we will see, this depends on the *metric space* in which the vertices lie. By metric space, we mean a space on which a distance can be applied. This is probably best illustrated with an example. The most often used metric in the plane is the Euclidean metric, which specifies that for any two points a and b, $dist(a,b) = \sqrt{(a_x - b_x)^2 + (a_y - b_y)^2}$. We refer to the two-dimensional space on which the Euclidean metric is applied as the *Euclidean plane.*

We can generalize the Euclidean metric to create a family of distance metrics known as the *Minkowski metrics*, after the mathematician of the same name. For $m \geq 1$, we define the d-dimensional metric L_m as the metric such that for any two points a and b,

$$dist(a,b) = \sqrt[m]{|a_1 - b_1|^m + |a_2 - b_2|^m + \cdots + |a_d - b_d|^m}. \tag{1.1}$$

We refer to L_m as the m^{th}-order Minkowski metric. Note that the Euclidean metric is a special case, the 2^{nd}-order Minkowski metric.

Let's examine the infinite-order Minkowski metric. In Equation 1.1, what happens when m approaches infinity? The largest difference in co-ordinate values dominates the expression, rendering all others negligible. Therefore, the distance between a and b in L_∞ is $\max\{|a_1 - b_1|, |a_2 - b_2|, \ldots, |a_d - b_d|\}$.

What does this have to do with the sphere of influence graph? The metric determines the shape of the circles drawn. We know from elementary school that in the Euclidean plane, we draw a circle with radius r about the origin with the equation $x^2 + y^2 = r^2$. The equation is true for all points distance r from the origin. However, if we are in the plane with the L_∞ metric, the equation is much different. Here, all points distance r from the origin satisfy the equation $\max\{x, y\} = r$. As we see in Figure 1.7, the circles are quite different! Spheres in the L_∞ metric are squares in two dimensions, cubes in three dimensions, and hypercubes in higher dimensions. It is important to realize that these spheres therefore have corners, a fact which we will exploit throughout Chapter 2.

Since the sphere of influence graph is based on circles, the resulting graphs are different as well. I would have also liked to demonstrate SIGs in more than two dimensions, but the limitations of paper prove too great.

In the next chapter, we present upper and lower bounds for the maximum size of the SIG in d-dimensional infinite-order Minkowski space (M_∞^d). In Chapter 3, we discuss a brief history of the problem and prove a new upper bound of $15n$ edges for the size of a SIG of n vertices in the Euclidean plane. Chapter 4 concludes the thesis and mentions a few open problems.

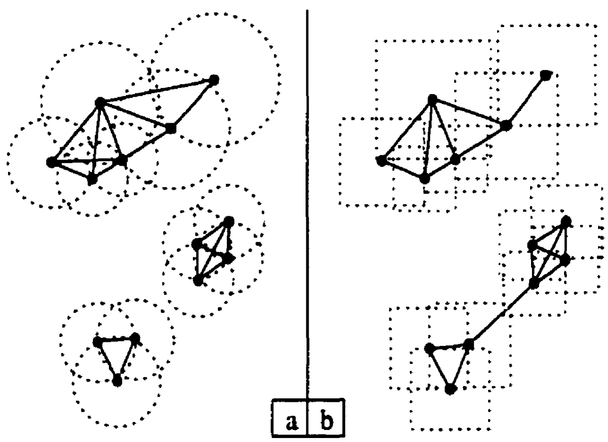

Figure 1.7: A set of points and its SIG in (a) the Euclidean plane and in (b) M_∞^2.

Chapter 2

Sphere of Influence Graphs in M_∞^d

In this chapter, we present upper and lower bounds for the maximum number of edges of the sphere of influence graph (SIG) in the d-dimensional infinite-order Minkowski space (M_∞^d). Section 2.1 contains a proof that no SIG of n vertices in M_∞^d has $(2^{2d-1} - 2^{d-1})n$ edges or more. In Section 2.2, we construct a lattice in M_∞^d whose number of SIG edges asymptotically approaches $(2^{d+2} - 3d - 4)n/9$ as n, the number of vertices, increases.

2.1 An upper bound on the number of edges

In this section, we demonstrate that no M_∞^d-SIG contains $(2^{2d-1} - 2^{d-1})n$ edges or more.

We start our proof by demonstrating that each intersection of spheres of influence contains at least two corners of the spheres. Although our spheres are open, we first prove the case for closed balls since the proof is easier.

Lemma 2.1 *Let X and Y be two intersecting* closed *balls in M_∞^d. Then the number of corners of X plus the number of corners of Y contained in $X \cap Y$ is at least two.*

Proof. Without loss of generality, let X be the smaller of the two balls. First we show that at least one corner of X is contained in Y. Because Y cannot be fully contained in X, at least part of one facet of X is contained in Y. Consider the cross-section of Y along the hyperplane of that facet. Since X is smaller than Y, that cross-section cannot fit inside the facet itself. Therefore there is some corner of X on that facet inside Y. Call that corner c.

Now suppose that c is the only corner of X inside Y; otherwise, our lemma is proven. This also implies that c is the only corner of $X \cap Y$ which is a corner of X. Note that the polytope $X \cap Y$ is an isothetic (axis-parallel) polytope, so let c' be the corner of $X \cap Y$ opposite from c, i.e., the corner which does not share a facet with c.

If X is to have only one corner inside $X \cap Y$, c' cannot be on the boundary of X. Suppose it were. Since c' shares no facets with c, two parallel supporting hyperplanes of X are distance cc' apart. Therefore, the diameter of X, and thus the length of an edge of X, is cc'. Since the length of some edge of $X \cap Y$ is also cc', an entire edge of X, and thus two vertices of X, belong to $X \cap Y$.

If c' is a corner of $X \cap Y$ not on the boundary of X, then c' is an intersection of facets only on the boundary of Y. Therefore, c' is a corner of Y. This completes the proof.

□

Lemma 2.1 discussed corners of *closed* balls; we now prove a corresponding lemma for *open* balls.

Since spheres of influence are open balls, their corners are not inside. This poses a problem, since in proving an upper bound for the size of the M_∞^d-SIG, we will require that we examine intersections at points contained inside the balls. To compensate, we examine points inside the balls very close to the corners. We refer to these points as ε-corners and define them as follows.

Let S be the spheres of influence of some point set. Let ε be one-half the smallest width of all polytopes determined by pairwise intersections of balls in S. The *width* of a polytope is the smallest distance between any pair of parallel hyperplanes which support the polytope. For each corner c of a ball in S, we define an *ε-corner* as $c + \vec{v}_{c\to o}$ where $\vec{v}_{c\to o}$ is the vector of length ε in the direction from the corner to the centre of the ball.

Lemma 2.2 *Let X and Y be two intersecting open balls in M_∞^d. Then the number of ε-corners of X plus the number of ε-corners of Y contained in $X \cap Y$ is at least two.*

Proof. If two open balls X and Y intersect, then their closures $\overline{X}$ and $\overline{Y}$ also intersect. By Lemma 2.1, two corners are contained in $\overline{X} \cap \overline{Y}$. Since X and Y are open, $\overline{X} \cap \overline{Y}$ has a non-zero width. Therefore, the ε-corners corresponding to the two corners of $\overline{X}$ and $\overline{Y}$ are in $X \cap Y$.

□

Now that we've shown that each intersection contains at least two ε -corners, we limit the number of intersections in which each ε -corner appears.

Lemma 2.3 *Let p be a point in M_∞^d, and S be a collection of open balls that do not contain each other's centres. Let Q be a closed orthant of M_∞^d whose corner lies at p. Then there is at most one ball in S which contains p and whose centre lies in Q.*

Proof. We prove this lemma by contradiction. Without loss of generality, let p be the origin and Q be the orthant determined by $x_i \geq 0$ for all $1 \leq i \leq d$. Suppose there are two balls, s_a and s_b, which contain p and whose centres, a and b, lie in Q. Let m_a be the distance from a to the origin, and let m_b be the distance from b to the origin. Without loss of generality, assume that $m_b \geq m_a$. A two-dimensional example is illustrated in Figure 2.1.

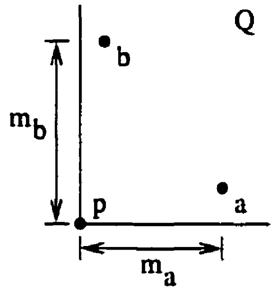

Figure 2.1: Lemma 2.3.

Since p is at the origin, m_b is also the largest co-ordinate value of b. Then the point $(\frac{m_b}{2}, \frac{m_b}{2}, \ldots, \frac{m_b}{2})$ is at most distance $\frac{m_b}{2}$ from b. Since s_b contains p, the radius of s_b is greater than m_b. Thus s_b must contain the closed ball centered at $(\frac{m_b}{2}, \frac{m_b}{2}, \ldots, \frac{m_b}{2})$ with radius $\frac{m_b}{2}$. Call this ball Z.

Since Z is centered at $(\frac{m_b}{2}, \frac{m_b}{2}, \ldots, \frac{m_b}{2})$ with radius $\frac{m_b}{2}$, Z contains the set $\{x : \textit{for all } 1 \le i \le d,\ 0 \le x_i \le m_b\}$. Because Z is contained in a sphere of influence, a cannot lie inside Z. Therefore there is one co-ordinate value of a that is greater than m_b. Thus $m_a > m_b$, and we have a contradiction.

□

Since every point in M_∞^d can be used to define a set of 2^d orthants, we can draw the following corollary.

Corollary 2.4 *Let S be a collection of open balls in M_∞^d such that no ball in S contains the centre of any other. Then no point in M_∞^d is contained in more than 2^d balls of S.*

We end this section with an upper bound on the size of a M_∞^d-SIG.

Theorem 2.5 *No sphere of influence graph of n vertices in M_∞^d has $(2^{2d-1} - 2^{d-1})n$ edges or more.*

Proof. Let S be the spheres of influence of some SIG G. Let C be the set of ε-corners of balls in S. Since S contains n balls, and each ball has 2^d ε-corners, C contains $2^d n$ points (some possibly duplicated).

Let X be a ball in S such that p is a ε-corner of X. By Corollary 2.4, each ε-corner in C is contained in no more than 2^d balls. Since $p \in X$, p is contained in at most $2^d - 1$ other balls of S. Therefore, p cannot be in more than $2^d - 1$ pairwise intersections of balls involving X.

C contains $2^d n$ ε-corners, each of which are involved in at most $2^d - 1$ intersections, and by Lemma 2.2, each intersection contains at least two ε-corners. Thus G contains at most

$$2^d n \times (2^d - 1) \div 2 = \frac{(2^{2d} - 2^d)n}{2} = (2^{2d-1} - 2^{d-1})n$$

edges. Since an ε-corner on the convex hull of C cannot be involved in $2^d - 1$ intersections, we can be a bit more precise in saying that G must have fewer than $(2^{2d-1} - 2^{d-1})n$ edges.

□

2.2 A lower bound on the maximum number of edges

In this section, we construct a lattice in M_∞^d space such that each vertex has $(2^{2d+3} - 6d - 8)/9$ sphere of influence neighbours. Readers who are not familiar with lattices or their generating bases may wish to read Appendix A on page 46 for a brief review of the subject.

The proof of the upper bound in Section 2.1 relies on the usage of corners in intersections, specifically that each corner can be used in at most $2^d - 1$ intersections. Thus, if we are to create a SIG with as many edges as possible, it seems reasonable

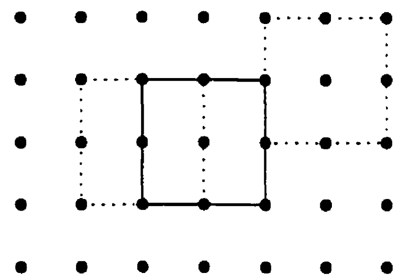

Figure 2.2: Two dimensional square lattice and spheres of influence.

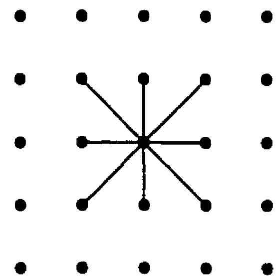

Figure 2.3: Two dimensional square lattice with SIG edges drawn from the centre vertex.

to seek a structure in which each intersection of spheres involves as few corners as possible, namely two (in light of Lemma 2.2). This is our aim in creating the lattice. We achieve a tight bound on the maximum size of the SIG in two dimensions, and trivially in one dimension. The gap between bounds is narrow in three dimensions but widens as the number of dimensions approaches infinity.

Figure 2.2 demonstrates the lattice generated by an orthogonal unit basis of the plane. If we compute the sphere of influence graph of the lattice shown in Figure 2.3, we see that the origin (and thus every point) has eight neighbours. The SIG resulting from the two-dimensional square lattice is shown in Figure 2.3. Note that there are some intersections which involve more than the minimum two corners. The intersections between north-south and east-west neighbours each involve four corners.

If we perturb the lattice a little bit, then these corners can be freed to intersect

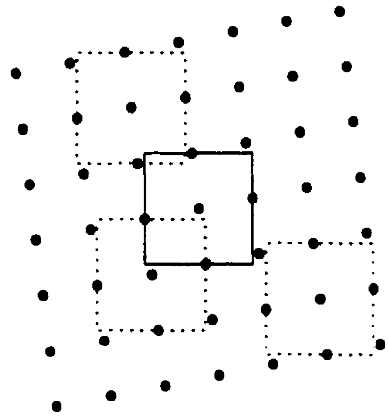

Figure 2.4: Subset of two dimensional tilted lattice and spheres of influence.

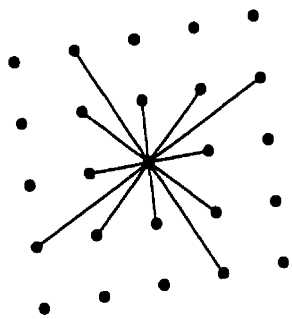

Figure 2.5: Subset of two dimensional tilted lattice with SIG edges drawn from the centre vertex.

other spheres. We apply rotations to the basis vectors of the square lattice, replacing $\{(1,0),(0,1)\}$ with the new unit[1] basis $\{(1,\delta_1),(-\delta_2,1)\}$, where δ_1,δ_2 are small positive integers. Note that it is not necessary that $\delta_1=\delta_2$; in fact, in higher dimensions it is essential that they are not equal! Thus the degrees of rotation applied to the individual basis vectors differ. This lattice is demonstrated in Figure 2.4.

As shown in Figure 2.5, these rotations provide us with four more neighbours per vertex than does the square lattice. This tilting is precisely the motivation for our *d*-

[1]The reader is reminded that in M_∞^2 the vector $(1,\delta)$ is a unit vector if $|\delta|\le 1$.

dimensional lattice that we present next, to which we refer as the *tilted d-dimensional lattice*, which will be abbreviated as T_d. (The lattice in Figure 2.4 would then be appropriately labelled T_2.)

2.2.1 A generating basis for a tilted d-dimensional lattice

Let $\widehat{u_j}$ be the unit vector for the j^{th} dimension. (The $\widehat{u_j}$'s are orthogonal.) We generate the lattice with the basis $\mathcal{B} = \{\vec{b_i} : 1 \leq i \leq d\}$ such that

$$\vec{b_i} = \widehat{u_i} + \delta_i \sum_{j<i} \widehat{u_j} - \delta_i \sum_{j>i} \widehat{u_j} \tag{2.1}$$

where each δ_i is a small positive real number, such that $\delta_1 \leq 1/4$, and $\delta_j \leq \frac{\delta_{j-1}}{4}$ for all $2 \leq j \leq d$. The basis $\mathcal{B}$ can also be visualized by

$$\begin{aligned}
\vec{b_1} &= (1, -\delta_1, -\delta_1, -\delta_1, -\delta_1 \ldots) \\
\vec{b_2} &= (\delta_2, 1, -\delta_2, -\delta_2, -\delta_2, \ldots) \\
\vec{b_3} &= (\delta_3, \delta_3, 1, -\delta_3, -\delta_3, \ldots) \\
\vec{b_4} &= (\delta_4, \delta_4, \delta_4, 1, -\delta_4, \ldots) \\
&\vdots \\
\vec{b_d} &= (\delta_d, \delta_d, \delta_d, \delta_d, \ldots, \delta_d, 1).
\end{aligned}$$

The δ_i's have been carefully chosen to facilitate the counting of SIG edges; this will be evident in Subsection 2.2.2. We end this subsection with a lemma showing the relationship between any term δ_i and the sum of its successive δ's.

Lemma 2.6 *For any i,*

$$\delta_i > \sum_{j>i} 2\delta_j.$$

Proof. Since $\delta_j \leq \frac{\delta_i}{4^{j-i}}$, the summation is less than the geometric series

$$\frac{2\delta_i}{4} + \frac{2\delta_i}{4^2} + \frac{2\delta_i}{4^3} + \cdots$$

which converges to $2\delta_i/3$.

□

2.2.2 Counting the SIG edges of T_d

We now come to the task of counting the number of sphere of influence edges in a tilted d-dimensional lattice.

For the remainder of this section, we will often refer to a lattice point x as a vector in the co-ordinate system determined by our basis $\mathcal{B}$. We refer to this vector as the T-vector of the point x. For example, if a lattice point in T_4 is determined by

$$3\vec{b_1} - 4\vec{b_2} + \vec{b_4}$$

then its T-vector would be $(3, -4, 0, 1)$. In an axis-parallel unit basis of M_∞^4, this vector has the co-ordinates

$$(3 + 4\delta_2 + \delta_4, -4 - 3\delta_1 + \delta_4, -3\delta_1 + 4\delta_2 + \delta_4, 1 - 3\delta_1 + 4\delta_2).$$

We refer to this vector as the $\hat{u}$-vector. Simply put, the T-vector of a point x is the vector of x in the co-ordinate system defined by $\mathcal{B}$, and the $\hat{u}$-vector of x is the vector of x in the co-ordinate system defined by axis-parallel unit vectors in M_∞^d.

Lemma 2.7 *The length of any vector in M_∞^d is greater than 2/3 the absolute value of the largest co-ordinate value in its corresponding T-vector.*

Proof. Let $\vec{t}$ be the T-vector of some point in M_∞^d, and let $\vec{w}$ be its $\hat{u}$-vector. Suppose the largest absolute value m of $\vec{t}$ occurs at the i^{th} co-ordinate. Then the i^{th} value of $\vec{w}$ is m plus a sum of δ_j's. To be precise,

$$w_i = t_i + \sum_{j<i} t_j\delta_j - \sum_{j>i} t_j\delta_j.$$

The absolute value of w_i is bounded by the expression

$$|w_i| \geq m - \textstyle\sum_{j\neq i} |t_j\delta_j|$$

$$|w_i| \geq m - \textstyle\sum_{j=1}^{d} |t_j\delta_j|$$

$$|w_i| \geq m - \textstyle\sum_{j=1}^{d} |m\delta_j|$$

$$|w_i| \geq m(1 - \textstyle\sum_{j=1}^{d} \delta_j).$$

Recall that the δ_j's are determined by $\delta_1 \leq 1/4$, and $\delta_j \leq \frac{\delta_{j-1}}{4}$ for all $2 \leq j \leq d$. The summation $\sum_{j=1}^{d} |\delta_j|$ is thus at most a finite sum of the sequence

$$\frac{1}{4} + \left(\frac{1}{4}\right)^2 + \left(\frac{1}{4}\right)^3 + \cdots$$

whose infinite sum converges to $1/3$. Since the finite sum is less than $1/3$, $|w_i| > 2m/3$.

$\square$

Lemma 2.8 *The sphere of influence of every vertex in T_d has radius* 1.

Proof. We must show that for any vertex p, its nearest neighbour is distance 1 away. This neighbour is no more than distance 1 away, since there exists a vertex distance 1 from p, namely $p + \vec{b_1}$. To complete the proof, we must demonstrate that there exists no point closer than 1.

Let x be a vertex of T_d distinct from p. Let $\vec{t}$ be the T-vector of $x - p$, which is non-zero. If any co-ordinate value of $\vec{t}$ is not in {-1, 0, 1}, then by Lemma 2.7 the length of $x - p$ is at least $4/3$. Thus we need only focus on T-vectors $\vec{t}$ whose co-ordinate values all lie in the set {-1, 0, 1}.

If $\vec{t}$ has only one non-zero value t_i, then the length of $x - p$ is $|t_i|$, which is 1.

If $\vec{t}$ has more than one non-zero value, we focus on its first two non-zero values, t_a and t_b $(a < b)$. Let $\vec{w}$ be the $\hat{u}$-vector of $x - p$. Then w_a and w_b can be described by

$$w_a = t_a + t_b\delta_b + \textstyle\sum_{j>b} t_j\delta_j;$$
$$w_b = t_b - t_a\hat{\delta}_a + \textstyle\sum_{j>b} t_j\hat{\delta}_j.$$

By Lemma 2.6,

$$|t_a\delta_a| > |t_b\delta_b| > \sum_{j>b} |t_j\hat{\delta}_j|.$$

Therefore, if t_a and t_b have the same sign, then $|w_a| > 1$ regardless of the summation. if t_a and t_b have different signs, then $|w_b| > 1$. Thus the length of $x - p$ is at least 1.

□

Because every sphere of influence in the lattice has radius 1, two spheres intersect if the distance between their centres is less than 2.

Corollary 2.9 *For any two distinct points $x, y \in T_d$, $\{x, y\}$ is a SIG edge if and only if $dist(x, y) < 2$.*

Lemma 2.10 *Let x and y be two lattice points of T_d. Let $\vec{t}$ be the T-vector of $y - x$. Then $dist(x, y) < 2$ if and only if one of the following two cases is true:*

- *Case 1 (all three of the following statements are true)*
  - *The first non-zero co-ordinate value of $\vec{t}$ is either -1 or $+1$.*
  - *If the first non-zero value of $\vec{t}$ is -1, then all co-ordinate values of $\vec{t}$ are in the set $\{-2, -1, 0, +1\}$.*
  - *If the first non-zero value of $\vec{t}$ is $+1$, then all co-ordinate values of $\vec{t}$ are in the set $\{-1, 0, +1, +2\}$.*

- *Case 2 (all four of the following statements are true)*
  - *The first non-zero co-ordinate value of $\vec{t}$ is either -2 or $+2$.*
  - *There exists a second non-zero co-ordinate value of $\vec{t}$.*
  - *If the first non-zero value of $\vec{t}$ is -2, then the second non-zero value of $\vec{t}$ is $+1$, and all co-ordinate values of $\vec{t}$ are in the set $\{-2, -1, 0, +1\}$.*
  - *If the first non-zero value of $\vec{t}$ is $+2$, then the second non-zero value of $\vec{t}$ is -1, and all co-ordinate values of $\vec{t}$ are in the set $\{-1, 0, +1, +2\}$.*

Proof. Without loss of generality, consider x to be the origin, and y to be a vertex of T_d distinct from x. Thus $\vec{t}$ is the T-vector of $\vec{y}$. Let $\vec{w}$ be the $\hat{u}$-vector of $\vec{y}$. Recall that in the metric space M_∞^d, y will be distance 2 or greater from x if one of the co-ordinate values of $\vec{w}$ is at least 2.

Consider $\vec{w}$ expressed as

$$(\alpha_1 + \gamma_1, \alpha_2 + \gamma_2, \ldots, \alpha_d + \gamma_d)$$

where $(\alpha_1, \alpha_2, \ldots, \alpha_d)$ is $\vec{t}$, and where the γ_i's are the sums of the δ's from Equation 2.1. Thus, the γ_i's can be expressed as

$$\gamma_i = \sum_{j>i} \alpha_j \delta_j - \sum_{j<i} \alpha_j \delta_j. \tag{2.2}$$

Now we show that $\vec{t}$ satisfies one of the two cases above. From Lemma 2.7, we know that if y is within distance 2 of x, then all co-ordinate values of $\vec{t}$ (and thus all the α's) are in the set $\{-2, -1, 0, +1, +2\}$. Thus, we assume that the α's are in this set.

Let α_a be the first non-zero value of $\vec{t}$, and let α_b be the second non-zero value, if it exists. Then by Lemma 2.6, we can see by examining Equation 2.2 that for all $i > a$, the sign of γ_i is dominated by the sign of $-\alpha_a$. The sign of γ_a is dominated by the sign of α_b, if it exists.

Due to the geometric progression of the δ's, $-1 < \gamma_i < 1$ for all i. Since $\alpha_i \in \{-2, -1, 0, +1, +2\}$, y is not within distance 2 of x if and only if there exists some k such that either $\alpha_k = +2$ and γ_k is nonnegative, or such that $\alpha_k = -2$ and γ_k is nonpositive. If $\alpha_a = +1$, then for $i > a$, all γ_i are negative, and therefore $\alpha_i \in \{-1, 0, +1, +2\}$. Likewise, if $\alpha_a = -1$, then for all $i > a$, all γ_i are positive, and therefore $\alpha_i \in \{-2, -1, 0, +1\}$. This proves case 1.

If $\alpha_a = +2$, again we draw the conclusion that for all $i > a$, $\alpha_i \in \{-1, 0, +1, +2\}$. However, unlike case 1, γ_a becomes important. In order that y be within distance 2 of x, γ_a must be negative. This implies the existence of the second non-zero value α_b, which must be negative to achieve the desired sign of γ_a. Since γ_b is negative, the only possibility for α_b is -1. A symmetric argument covers the case for $\alpha_a = -2$. This proves case 2.

Since there are no other possibilities for α_a, the proof is complete.

□

Theorem 2.11 *In the infinite lattice T_d, each vertex has $(2^{2d+3} - 6d - 8)/9$ sphere of influence neighbours.*

Proof. As shown by Corollary 2.9, this amounts to counting the number of vertices that are within distance 2 of the origin. Thus we simply have to count the vectors that satisfy the conditions of Lemma 2.10. Let $\vec{t}$ be such a vector such that the first co-ordinate value of $+1$ or -1 occurs at the i^{th} position. Then $\vec{t}$ follows either case 1 or case 2 as outlined in Lemma 2.10.

If $\vec{t}$ follows case 1, then all co-ordinate values of t preceding t_i are 0, and all succeeding values have four possibilities, -1, 0, $+1$, and either -2 or $+2$, depending on t_i. Thus for a d-dimensional vector following case 1, there are two choices for the i^{th} value, and after that four choices for each of the $i + 1^{th}$ through d^{th} values. Therefore, there are $2(4^{d-i}) = 2^{2d-2i+1}$ choices for vectors following case 1 whose first non-zero value occurs at the i^{th} position.

If $\vec{t}$ follows case 2, then t_i must be the second non-zero value of $\vec{t}$. Then exactly one value preceding t_i, say t_k, is $-2t_i$, and all values after t_i have four possibilities, -1, 0, $+1$, and t_k. For a d-dimensional vector following case 2, there are two choices for the i^{th} value, each of which imply a single value for the k^{th} value. However, there are $i-1$ choices for k, since $1 \leq k \leq i-1$. After these values have been decided, there are four choices for each of the $i+1^{st}$ through d^{th} values. Therefore, there are $(i-1)(2)(4^{d-i}) = (i-1)2^{2d-2i+1}$ choices for vectors following case 2 whose first non-zero value occurs at the i^{th} position.

Summing over cases 1 and 2 yields $2^{2d-2i+1}i$ possible vectors for any given i. Summing this over all possible values of i from 1 to d yields

$$\sum_{i=1}^{d} 2^{2d-2i+1}i = 2^{2d-1} + 2^{2d-3}(2) + 2^{2d-5}(3) + \cdots + 2^3(d-1) + 2(d) \tag{2.3}$$

vectors. To facilitate the computation, we rearrange the terms to yield a sum of finite geometric series. Summation 2.3 can be expanded as

$$\left.\begin{array}{r} 2 \\ +2^3 + 2 \\ \vdots \\ +2^{2d-5} + \cdots + 2^3 + 2 \\ +2^{2d-3} + 2^{2d-5} + \cdots + 2^3 + 2 \\ +2^{2d-1} + 2^{2d-3} + 2^{2d-5} + \cdots + 2^3 + 2. \end{array}\right\} d \text{ lines}$$

For $1 \leq k \leq d$, the k^{th} line is the geometric series

$$2^{2k-1} + 2^{2k-3} + \cdots + 2^3 + 2 = \frac{2^{2k-1} - \frac{1}{2}}{1 - \frac{1}{4}} = \frac{2^{2k+1} - 2}{3}.$$

Summing all of the lines reveals the total number of possible vectors;

$$\sum_{k=1}^{d} \frac{2^{2k+1} - 2}{3} = \frac{-2d}{3} + \sum_{k=1}^{d} \frac{2^{2k+1}}{3}.$$

This latest summation is another geometric series, so

$$\frac{-2d}{3}+\sum_{k=1}^{d}\frac{2^{2k+1}}{3}=\frac{-2d}{3}+\frac{2^{2d+1}-2}{(3)(1-\frac{1}{4})}=\frac{-2d}{3}+\frac{2^{2d+3}-8}{9}=\frac{2^{2d+3}-6d-8}{9}.$$

Therefore the origin has $(2^{2d+3}-6d-8)/9$ sphere of influence neighbours. Since a lattice is invariant to translations by integral multiples of its basis vectors, every vertex in T_d has $(2^{2d+3}-6d-8)/9$ sphere of influence neighbours.

□

Corollary 2.12 *There exist SIGs of n vertices in M_∞^d whose numbers of edges asymptotically approach* $(2^{2d+2}-3d-4)n/9$.

2.3 Conclusion

We end Chapter 2 with Table 2.1, which presents the bounds on the maximum size of the sphere of influence graph for a few selected dimensions.

Dimension (d)	Upper Bound	T_d	Gap between bounds
1	n	n	tight
2	$6n$	$6n$	tight
3	$28n$	$27n$	n
4	$120n$	$112n$	$8n$
5	$496n$	$453n$	$43n$
6	$2016n$	$1818n$	$198n$
$\lim_{d\to\infty}$			(1/9)Upper bound

Table 2.1: Selected bounds on the maximum number of edges in a M_∞^d-SIG.

Chapter 3

Sphere of Influence Graphs in the Euclidean Plane

In this chapter, we begin by giving a brief history of the problem, "What is the minimum value of c such that a sphere of influence of n vertices in the Euclidean plane (E-SIG) has at most cn edges?" In Section 3.2 we present a new upper bound of 15.

3.1 A brief history of the problem

Five years after the sphere of influence graph was introduced by Toussaint [Tou80], the question, "Does there exist a constant c such that a E-SIG has at most cn edges?" was solved by David Avis and Joe Horton [AH85], who provided the constant $c = 29$. They proved that given a sphere of influence graph $G(V)$ on a point set V, the vertex x_1 that has the smallest sphere of influence has at most 29 incoming edges. Any edge of $G(V)$ not touching x_1 is an edge of $G(V \setminus \{x_1\})$ since removing x_1 can only increase the radii of the spheres. That $G(V)$ contains at most $29n$ edges now follows by induction on the cardinality of V.

It was later realized that the theorem by Avis and Horton had been proven in a different form forty years earlier. In 1945, Abram Besicovitch required (and proved) the following lemma [Bes45]:

Lemma 3.1 (Besicovitch, 1945) *Given a set Γ of coplanar circles, the center of no one of them being in the interior of another, and U the circle (or a circle) of Γ whose radius does not exceed the radius of any other circle of Γ, then the number of circles meeting U does not exceed 21.*

The number 21 was improved to its lowest possible at 18 by E. R. Reifenberg in 1948 [Rei48] and independently by Paul Bateman and Paul Erdös in 1951 [BE51]. Since planar Euclidean spheres of influence are a collection of circles such that no interior of any circle contains the centre of any other, Lemma 3.1 can be reworded to apply directly to sphere of influence graphs. Thus by induction we can show that no sphere of influence graph of n vertices contains more than $18n$ edges. We can also make a statement concerning a similar graph, the *closed sphere of influence graph*, in which the spheres of influence are closed balls rather than open. Therefore we draw an edge between two vertices if their spheres intersect, whether or not the intersection is proper (has non-zero area). Here too the upper bound on the maximum size is $18n$.

We can reduce this bound to $17.5n$ with a simple realization. Let x_1 be the vertex with the smallest sphere of influence, of radius r_1. This sphere has radius r_1 because the nearest neighbour of x_1, say x_2 is distance r_1 away. Since r_1 is the smallest distance between any two vertices, then x_1 is also the nearest neighbour of x_2. Thus r_1 and r_2 (the radius of the sphere of influence of x_2) are both the smallest radii over all spheres, so x_1 and x_2 each have at most 18 neighbours. One edge is shared by x_1 and x_2, so the two vertices have in total 35 edges. Performing the induction on two vertices at a time instead of one yields a bound of $35n/2$ edges, or $17.5n$. This bound is attributed to Katchalski.

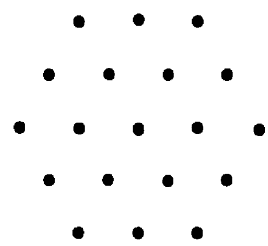

Figure 3.1: Subset of the hexagonal lattice.

Where is this upper bound headed? The aim, of course, is to find the optimal constant, joining the upper and lower bounds. For an idea of the tight bound, we consider the closed sphere of influence graph. We see that the hexagonal lattice has 18 neighbours per vertex, for $9n$ edges in total. In Figure 3.1, the centre vertex is a closed SIG neighbour of all the other drawn vertices. David Avis conjectures that the hexagonal lattice is optimal in that $9n$ is the most number of edges possible for a closed E-SIG. Since the open E-SIG is a subset of the closed E-SIG, the conjecture implies that the tight bound for the open E-SIG is no more than $9n$.

The problem was generalized to Euclidean spaces of arbitrary dimension in 1993 by Leonidas Guibas, János Pach, and Micha Sharir [GPS94]. They define the k^{th} *sphere of influence* ($k \geq 1$) of a point x as the open ball centered at x with a radius equal to the distance between x and its k^{th} nearest neighbour. The k^{th} *sphere of influence graph* is then the graph where two points are connected if their k^{th} spheres of influence intersect. A 2^{nd} sphere of influence graph is illustrated in Figure 3.2. The authors prove that for any d-dimensional Euclidean space, there exists a constant c_d (which depends exponentially on d) such that the k^{th} sphere of influence graph of n vertices contains at most $c_d kn$ edges.

Rex Dwyer also posed a related problem [Dwy93], "What is the expected size of the sphere of influence graph?" He proves that if n points are uniformly distributed in the d-dimensional unit ball, the expected number of edges lies between $(0.162)2^d n$

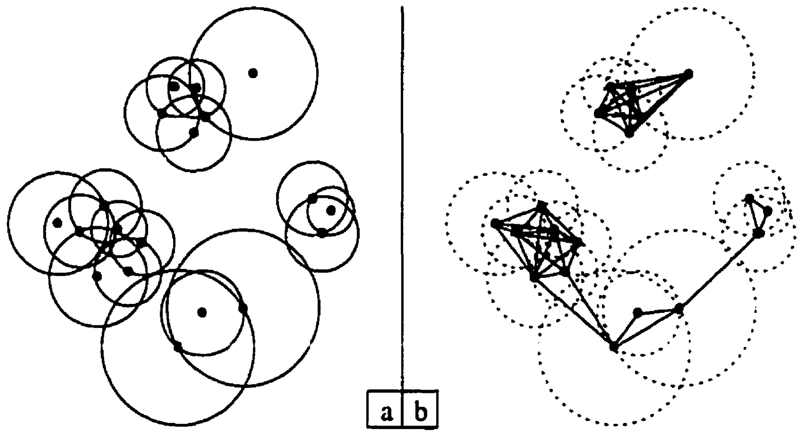

Figure 3.2: A set of points, (a) its 2^{nd} spheres of influence, and (b) its 2^{nd} E-SIG.

and $(0.667)2^d n$, for $d \geq 7$ and $n \gg d$.

While these related problems are interesting and worth noting, the remainder of this chapter concerns itself with the size of the sphere of influence graph in the Euclidean plane. In the next section we improve the upper bound on the size of the E-SIG to $15n$.

3.2 An upper bound of $15n$

In this section, we prove the following theorem, which is the main result of Chapter 3.

Theorem 3.2 *No open or closed sphere of influence graph of n vertices in the Euclidean plane contains more than $15n$ edges.*

To facilitate our proof, we assign *weights*, or numerical values, to the edges of the E-SIG as follows. First, we replace each undirected edge $\{a, b\}$ with two directed edges, (a, b) and (b, a). Let the radii of the spheres of influence of a and b be r_a and r_b, respectively. Then (a, b) is given a weight of 1 if $r_a \leq 2r_b/3$, a weight of $1/2$ if

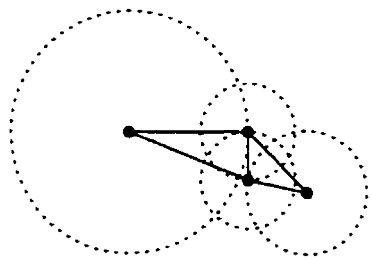

Figure 3.3: A sphere of influence graph in the Euclidean plane.

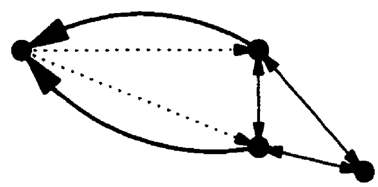

Figure 3.4: A weighted sphere of influence graph.

$2r_b/3 < r_a < 3r_b/2$, and a weight of 0 otherwise.[1] We refer to this graph as the *weighted sphere of influence graph*, or WSIG. A E-SIG is shown in Figure 3.3 and its corresponding WSIG in Figure 3.4. The thick lines represent edges of weight 1; the thin lines each represent a pair of edges of weight 1/2, and the dotted lines represent edges of weight 0.

Our goal is to utilize the WSIG in determining a new upper bound for the E-SIG.

Lemma 3.3 *On any point set V, the total weight of all edges in the WSIG of V is equal to the number of edges in the SIG of V.*

Proof. Each edge in the SIG corresponds to a pair of edges in the WSIG. Either both edges have weight 1/2, or one has weight 1 and the other 0. Thus each SIG edge corresponds to two WSIG edges whose weights add up to 1.

□

[1]It is remarked that the values 2/3 and 3/2 have been chosen to produce the best results as determined through trial and error.

Lemma 3.3 implies that if we can prove that no WSIG of n vertices has edges whose total weight is greater than $15n$, then we have also proven Theorem 3.2. This is exactly the method behind our proof, and we begin with the following theorem.

Theorem 3.4 *There exists no node in the WSIG for which the weights of outgoing edges sum to greater than 15.*

We will prove this theorem by demonstrating that it follows from Theorem 3.5 and then by proving the latter, which discusses fitting points into annuli.

Theorem 3.5 (modified from Reifenberg and Bateman-Erdös, 1948/1951) *Let the term* admissible point of weight $1/2$ *refer to a point p in the annulus $1 \leq \rho \leq 5/3$ such that no other admissible point is within distance $2/3$ of p. Let the term* admissible point of weight 1 *refer to a point q in the annulus $1.5 \leq \rho \leq 2.5$ such that*

- *no other admissible point of weight 1 is within distance 1.5 of q,*
- *no admissible point of weight $1/2$* not *on the circle $\rho = 5/3$ is within distance 1.5 of q, and*
- *for each admissible point of weight $1/2$ which has polar co-ordinates $(5/3, \theta)$, there exists a point in space (r, θ) where $5/3 \leq r \leq 2.5$ such that (r, θ) is at least distance 1.5 from q.*

Then it is impossible to fit any combination of admissible points in the annulus $1 \leq \rho \leq 2.5$ such that their total weights sum to a value greater than 15.

We delay the proofs of the last two theorems for now and instead prove that Theorem 3.5 implies Theorem 3.4. We first require the following lemma.

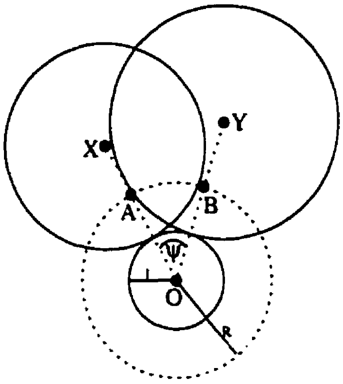

Figure 3.5: Lemma 3.6.

Lemma 3.6 *In polar co-ordinates, let $X = (x, \theta_x)$ and $Y = (y, \theta_y)$ be the centres of two circles that do not contain each other's centres but that both intersect $\rho = 1$. Furthermore, we impose the condition that X and Y lie outside the disk $\rho \leq R$, for some $R > 1$. Then the points $A = (R, \theta_x)$ and $B = (R, \theta_y)$ are at least distance $R - 1$ apart.*

Proof. As above, let $x = OX$, $y = OY$, and $\psi = m\angle XOY$, as shown in Figure 3.5. Then $dist(X, Y)^2 \geq \max\{(x-1)^2, (y-1)^2\}$, so

$$x^2 + y^2 - 2xy\cos\psi \geq \max\{(x-1)^2, (y-1)^2\}.$$

Supposing $x \leq y$ yields

$$\begin{aligned}\cos\psi \leq \frac{x^2 + y^2 - (y-1)^2}{2xy} &= \frac{1}{x} + \frac{x^2 - 1}{2xy} \\ &\leq \frac{1}{x} + \frac{x^2 - 1}{2x^2}\end{aligned}$$

Let this upper bound be $f(x)$. Differentiating with respect to x gives us

$$\frac{df}{dx} = -\frac{1}{x^2} + \frac{1}{4x^3}$$

which is negative for any $x > 1$. Thus $f(x)$ is strictly decreasing as x grows. Since the largest possible value for $\cos\psi$ is $f(R)$, the smallest value possible for ψ is $\arccos f(R)$. Then $m\angle AOB$ is also at least $\arccos f(R)$. Therefore,

$$\begin{aligned} dist(A,B)^2 &\geq 2R^2 - 2R^2 f(R) \\ &\geq 2R^2\left(\frac{1}{R} + \frac{R^2-1}{2R^2}\right) \\ &\geq R^2 - 2R + 1 \\ &\geq (R-1)^2. \end{aligned}$$

Thus $dist(A,B) \geq R - 1$. This proves the lemma.

□

We are now ready to prove the following theorem.

Theorem 3.7 *Theorem 3.5 implies Theorem 3.4.*

Proof (generalized from a proof by Bateman and Erdős). Let O be some node in the WSIG with at least one outgoing edge of non-zero weight. Without loss of generality, assume that O is at the origin and that the sphere of influence of O has radius 1. Thus we have a set Δ of circles of radius at least 2/3 which intersect the circle $\rho = 1$ such that the centre of no circle is contained in any other. Also, since the sphere of influence of O has radius 1, no circle in Δ is centered inside $\rho < 1$.

It suffices to show that we can construct a set Δ^* of admissible points where each circle in Δ of radius more than 2/3 but less than 1.5 corresponds to a point in Δ^* with weight 1/2, and where each circle in Δ of radius 1.5 or greater corresponds to a point in Δ^* with weight 1. Furthermore, we demand that both correspondences be *bijective*, meaning that every circle in Δ corresponds uniquely to a point in Δ^* and vice versa.

First, let us choose the points of weight 1. We select a point of weight 1 corresponding to a circle C of radius 1.5 or greater in the following manner: if the centre of C lies inside $\rho \leq 2.5$, we select its centre; if not, we select the point lying on the circle $\rho = 2.5$ with the same amplitude. As a result, the circle of radius 1.5 centered at a point of weight 1 in Δ^* is contained in its corresponding circle in Δ, which contains no other centres. Thus it remains only to demonstrate that if two circles of Δ have centres X and Y outside $\rho \leq 2.5$, their corresponding points a and b are at least distance 1.5 apart. This is proven by Lemma 3.6 since a and b are on the circle $\rho = 2.5$.

Now we are left with the selection of points of weight 1/2, which we perform in a similar manner. We select a point of weight 1/2 corresponding to a circle C in the following manner: if the centre of C lies inside $\rho \leq 5/3$, we select its centre; if not, we select the point lying on the circle $\rho = 5/3$ with the same amplitude. By the same logic of the proof concerning points of weight 1 and by Lemma 3.6, all points in Δ^* of weight 1/2 are mutually at least distance 2/3 apart.

It remains to prove that every point of weight 1 is at least distance 1.5 from points of weight 1/2 inside $\rho < 5/3$. We must also show that for each admissible point p of weight 1 and every admissible point with polar co-ordinates $(5/3, \theta)$ of weight 1/2, there exists a point in space (r, θ) where $r \geq 5/3$ such that (r, θ) is least distance 1.5 from p. Since the circle of radius 1.5 around each point p of weight 1 is contained in the corresponding circle of p, no circle of weight 1/2 can be centered within distance 1.5 of p. Thus, the only possibility in which a point q of weight 1/2 is within distance 1.5 of p is if q corresponds to a circle centered elsewhere. Then q would have to be on the circle $\rho = 5/3$. However, the centre of the corresponding circle is at least distance 1.5 away, even if q is not. This centre lies somewhere on the line $5/3 \leq \rho \leq 2.5$ with the same amplitude as q, so there is some point on that line which lies at least distance 1.5 from p. This constraint is precisely the exception with regard to points of weight

1/2 lying on $\rho = 5/3$. This completes the proof of the lemma.

□

We now begin our proof of Theorem 3.5.

Lemma 3.8 (generalized from a lemma by Bateman-Erdös, 1951) *Label the origin as O. Let r, R, and τ be such that $0 < R - \tau \le r \le R$. Suppose that we have two points P and Q which lie in the annulus $r \le \rho \le R$ and which have mutual distance τ. Then the minimum value $\Phi_\tau(r, R)$ of $m\angle POQ$ has the smaller of the two values*

$$\Phi_\tau(r, R) = \arccos \frac{(R/\tau)^2 + (r/\tau)^2 - 1}{2Rr/\tau}, \textit{ and}$$
$$\Phi_\tau(r, R) = \arccos(1 - \frac{1}{2(R/\tau)^2}) = 2 \arcsin \frac{\tau}{2R}.$$

Proof. It suffices to consider the case where $OQ = R$ and $PQ = \tau$. Let $OP = \rho$. Our problem can be reduced to finding the ρ which yields the minimum value of $m\angle POQ$. Let $f(\rho) = m\angle POQ = \arccos[(R)^2 + (\rho)^2 - \tau^2)/(2R\rho)]$ for ρ in the interval r $\le \rho \le$ R. If we differentiate, we see that $f(\rho)$ cannot have an interior minimum in this interval. Thus the minimum is the smaller of $f(r)$ and $f(R)$, which are the two values described in the lemma.

□

The following instances will be used in the course of the proof of Theorem 3.5.

$$\Phi_{2/3}(1, 1.15) > 33^\circ.5 \qquad \Phi_{2/3}(1, 1.2) > 32^\circ.2$$
$$\Phi_{2/3}(1, 1.3) > 29^\circ.7 \qquad \Phi_{2/3}(1, 1.4) > 26^\circ.0$$
$$\Phi_{2/3}(1.15, 1.4) > 29^\circ.7 \qquad \Phi_{2/3}(1.15, 5/3) > 17^\circ.5$$
$$\Phi_{2/3}(1.2, 1.4) > 27^\circ.5 \qquad \Phi_{2/3}(1.2, 5/3) > 19^\circ.3$$

$$\Phi_{2/3}(1.3, 5/3) > 21^\circ.8 \qquad \Phi_{2/3}(1.4, 5/3) > 23^\circ.0$$

$$\Phi_{1.5}(1, 1.6) > 55^\circ.9 \qquad \Phi_{1.5}(1.15, 2.5) > 22^\circ.2$$

$$\Phi_{1.5}(1.2, 2.5) > 24^\circ.9 \qquad \Phi_{1.5}(1.3, 2.5) > 28^\circ.9$$

$$\Phi_{1.5}(1.4, 2.5) > 31^\circ.6 \qquad \Phi_{1.5}(1.5, 2.5) > 33^\circ.5$$

$$\Phi_{1.5}(1.6, 2.5) > 34^\circ.9$$

For the remainder of this section, the method of proof is simple. We place a configurations of admissible points inside annuli, sort the points radially, and compute the minimum angles between each pair of radial neighbours. For example, the angle between an admissible point of weight $1/2$ in $1 \leq \rho \leq 2$ and another in $1.3 \leq \rho \leq 1.4$ is at least $\Phi_{2/3}(1, 1.4)$. Likewise, between two admissible points of weight 1 both in $1.5 \leq \rho \leq 2.5$, the angle is at least $\Phi_{1.5}(1.5, 2.5)$. In general, let a, b, c, and d be such that $a \leq b \leq d$ and $a \leq c \leq d$. Then the angle between an admissible point of weight $1/2$ in $a \leq \rho \leq b$ and another in $c \leq \rho \leq d$ is at least $\Phi_{2/3}(a, d)$. In the case of points of weight 1, the angle is at least $\Phi_{1.5}(a, d)$.

Note that in between points of weight 1 and $1/2$, the distance is 1.5, unless the point of weight $1/2$ lies on $\rho = 5/3$. In this special case, the point of weight 1 must be distance 1.5 from some point in $5/3 \leq \rho \leq 2.5$ with the same amplitude as the point of weight $1/2$. Thus, we must always include the interval $5/3 \leq \rho \leq 2.5$ in this case. As an example, the minimum angle between a point of weight 1 in $1.5 \leq \rho \leq 2$ and a point of weight $1/2$ in $1.5 \leq \rho \leq 1.6$ is $\Phi_{1.5}(1.5, 2)$, but the minimum angle between a point of weight 1 in $1.5 \leq \rho \leq 2$ and a point of weight $1/2$ in $1.5 \leq \rho \leq 5/3$ is $\Phi_{1.5}(1.5, 2.5)$.

It is worth mentioning that this method requires a little more complication if we allow $a < c < d < b$, but since this case will not occur in our proof of Theorem 3.5, we will avoid discussing it here. Additionally, there are some cases where this method will not produce a tight lower bound on the angle. When $d - (\tau^2/d) < a$, then the minimum angle between a point in $a \leq \rho \leq b$ and another in $c \leq \rho \leq d$ is greater

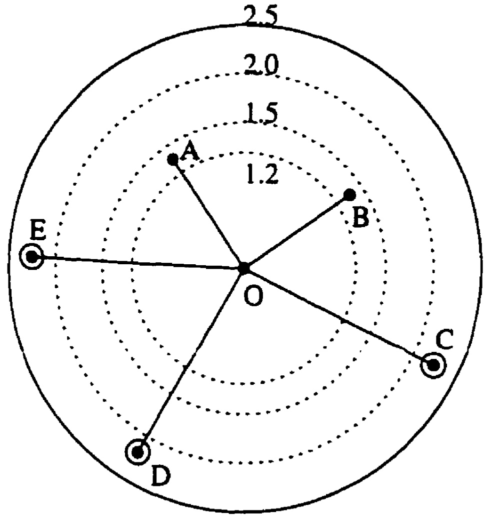

Figure 3.6: Sample configuration of admissible points.

than $\Phi_r(a,d)$. In all cases, $\Phi_r(a,d)$ still provides a valid lower bound, just not a tight bound. This does not affect our proof; it is only mentioned here for completeness.

A sample configuration of admissible points is demonstrated in Figure 3.6. The three encircled points have weight 1; the other two have weight 1/2. We see that $m\angle AOB \geq \Phi_{2/3}(1.2, 1.5)$, $m\angle BOC \geq \Phi_{1.5}(1.2, 2.5)$, and $m\angle COD \geq \Phi_{1.5}(2.0, 2.5)$.

The proof of each lemma below proceeds by assuming a possible configuration and then demonstrating that the sum of the subtended angles of radially consecutive points is greater than 360°. Since the configuration fits inside an annulus, we achieve a contradiction.

Lemma 3.9 *It is impossible to have 11 admissible points of weight 1.*

Proof. The angle between any two radially consecutive points (i.e., consecutive by amplitude) is at least $\Phi_{1.5}(1.5, 2.5)$. Since $\Phi_{1.5}(1.5, 2.5) > 33.5°$, our lemma follows from the fact that $11\Phi_{1.5}(1.5, 2.5) > 368°.5$.

□

Therefore the following six cases suffice to prove Theorem 3.5. We show that it is impossible to have

- Case 1: 10 admissible points of weight 1 and 11 of weight 1/2.
- Case 2: 9 admissible points of weight 1 and 13 of weight 1/2.
- Case 3: 8 admissible points of weight 1 and 15 of weight 1/2.
- Case 4: 7 admissible points of weight 1 and 17 of weight 1/2.
- Case 5: 6 admissible points of weight 1 and 19 of weight 1/2.
- Case 6: 26 admissible points of non-zero weight.

These cover all possibilities of having admissible points whose total weight exceeds 15.

The remainder of this chapter continues the proof of Theorem 3.5, which implies our main theorem that no E-SIG contains more than $15n$ edges. The reader who does not wish to see an exhaustive analysis of the above cases is advised to skip ahead to Chapter 4 on page 44.

To prove the individual cases, we require the following three lemmas, which will be used several times in the course of this chapter.

Lemma 3.10 *It is impossible to have 11 admissible points of weight 1/2 in the annulus $1 \le \rho \le 1.3$ such that 10 of them lie in $1 \le \rho \le 1.15$.*

Proof. Suppose we have 11 admissible points in $1 \le \rho \le 1.3$, 10 of which lie in $1 \le \rho \le 1.15$. Then 9 of the angles subtended at the origin by pairs of radially consecutive points are at least $\Phi_{2/3}(1, 1.15)$, and the other two are at least $\Phi_{2/3}(1, 1.3)$. However, $9\Phi_{2/3}(1, 1.15) + 2\Phi_{2/3}(1, 1.3) > 361^\circ.8$, so this is not possible.

□

Lemma 3.11 *It is impossible to have 12 admissible points of weight 1/2 in the annulus $1 \leq \rho \leq 1.4$ such that 10 of them lie in $1 \leq \rho \leq 1.2$.*

Proof. Suppose we have 12 such admissible points. Then either the 2 points in $1 \leq \rho \leq 1.4$ are each radially in between points in $1 \leq \rho \leq 1.2$, or the 2 points are radially consecutive. In the former case, the angles subtended are $8\Phi_{2/3}(1, 1.2) + 4\Phi_{2/3}(1, 1.4) > 361^\circ.6$. In the latter case, the angles subtended are $9\Phi_{2/3}(1, 1.2) + 2\Phi_{2/3}(1, 1.4) + \Phi_{2/3}(1.2, 1.4) > 369^\circ.3$. Thus we have a contradiction.

□

Lemma 3.12 *It is impossible to have 1 admissible point of weight 1 in the annulus $1.5 \leq \rho \leq 1.6$ and 9 admissible points of weight 1/2 in $1 \leq \rho \leq 1.2$.*

Proof. In this configuration, 9 of the angles subtended at the origin by pairs of radially consecutive points are at least $\Phi_{2/3}(1, 1.2)$, and the other two are at least $\Phi_{1.5}(1, 1.6)$. However, $9\Phi_{2/3}(1, 1.2) + 2\Phi_{1.5}(1, 1.6) > 369^\circ.4$, which proves the lemma.

□

We are now ready to proceed with the proofs of each of the six cases.

Lemma 3.13 (Case 1) *It is impossible to have 10 admissible points of weight 1 and 11 admissible points of weight 1/2.*

Sublemma 3.13.1 *It is impossible to have 10 admissible points of weight 1 in $1.5 \leq \rho \leq 2.5$ and 2 admissible points of weight 1/2 in $1.2 \leq \rho \leq 5/3$.*

Proof of sublemma. If the 2 points of weight 1/2 are consecutive, then the angles subtended are $9\Phi_{1.5}(1.5, 2.5) + 2\Phi_{1.5}(1.2, 2.5) + \Phi_{2/3}(1.2, 5/3) > 371^\circ.2$. If they are not consecutive, then the angles are $8\Phi_{1.5}(1.5, 2.5) + 4\Phi_{1.5}(1.2, 2.5) > 368^\circ.2$.

Since 10 admissible points must be in $1 \leq \rho \leq 1.2$, by Lemma 3.12 all admissible points of weight 1 are in $1.6 \leq \rho \leq 2.5$.

Sublemma 3.13.2 *It is impossible to have 10 admissible points of weight 1 in $1.6 \leq \rho \leq 2.5$ and 1 admissible point of weight 1/2 in $1.3 \leq \rho \leq 5/3$.*

Proof of sublemma. This follows from the fact that $9\Phi_{1.5}(1.6, 2.5) + 2\Phi_{1.5}(1.3, 2.5) > 371^\circ.9$, which demonstrates that this configuration is not achievable.

This proves that if we have 10 admissible points of weight 1 and 11 admissible points of weight 1/2, then all 11 points of weight 1/2 must lie in $1 \leq \rho \leq 1.3$. But by Lemma 3.10, 2 of the 11 points of weight 1/2 must lie in $1.15 \leq \rho \leq 1.3$. We now show that this is not possible.

Sublemma 3.13.3 *It is impossible to have 10 admissible points of weight 1 in $1.6 \leq \rho \leq 2.5$ and 2 admissible points of weight 1/2 in $1.15 \leq \rho \leq 1.3$.*

Proof of sublemma. If the 2 points of weight 1/2 are consecutive, then the angles subtended are $9\Phi_{1.5}(1.6, 2.5) + 2\Phi_{1.5}(1.15, 2.5) + \Phi_{2/3}(1.15, 1.3) > 388^\circ.2$. If they are not consecutive, then the angles are $8\Phi_{1.5}(1.6, 2.5) + 4\Phi_{1.5}(1.15, 2.5) > 368^\circ.0$.

This proves the lemma and therefore case 1.

□

Lemma 3.14 (Case 2) *It is impossible to have 9 admissible points of weight 1 and 13 admissible points of weight 1/2.*

Sublemma 3.14.1 *It is impossible to have 9 admissible points of weight 1 in $1.5 \leq \rho \leq 2.5$ and 4 admissible points of weight 1/2 in $1.2 \leq \rho \leq 5/3$.*

Proof of sublemma. No configuration of the 13 points exists such that the angles subtended sum to 360° or less. The smallest sum possible occurs when the 4 points of weight 1/2 are each between points of weight 1. The sum of angles in this case is $5\Phi_{1.5}(1.5, 2.5) + 8\Phi_{1.5}(1.2, 2.5) > 367^\circ.3$.

If there are at most 3 points of weight $1/2$ in $1.2 \leq \rho \leq 5/3$, then there must be at least 10 points of weight $1/2$ in $1 \leq \rho \leq 1/2$. By Lemma 3.11, this implies that there are not 12 points of weight $1/2$ within $1 \leq \rho \leq 1.4$. Therefore at least 2 of the 3 points of weight $1/2$ in $1.2 \leq \rho \leq 5/3$ are in $1.4 \leq \rho \leq 5/3$. Furthermore, by Lemma 3.12 all 9 points of weight 1 are in $1.6 \leq \rho \leq 2.5$. We now show that this is not possible.

Sublemma 3.14.2 *It is impossible to have 9 admissible points of weight 1 in $1.6 \leq \rho \leq 2.5$ and 2 admissible points of weight 1/2 in $1.4 \leq \rho \leq 5/3$.*

Proof of sublemma. No configuration of the 11 points exists such that the angles subtended sum to 360° or less. The smallest sum possible occurs when the 2 points of weight $1/2$ are consecutive. The sum of angles in this case is $8\Phi_{1.5}(1.6, 2.5) + 2\Phi_{1.5}(1.4, 2.5) + \Phi_{2/3}(1.4, 5/3) > 365^\circ.4$.

This proves the lemma and therefore case 2.

□

Lemma 3.15 (Case 3) *It is impossible to have 8 admissible points of weight 1 and 15 admissible points of weight 1/2.*

We follow very closely the proof of Case 2 (Lemma 3.14).

Sublemma 3.15.1 *It is impossible to have 8 admissible points of weight 1 in $1.5 \leq \rho \leq 2.5$ and 6 admissible points of weight 1/2 in $1.2 \leq \rho \leq 5/3$.*

Proof of sublemma. No configuration of the 14 points exists such that the angles subtended sum to 360° or less. The smallest value possible occurs when the 6 points of weight $1/2$ are each between points of weight 1. The sum of angles in this case is $2\Phi_{1.5}(1.5, 2.5) + 12\Phi_{1.5}(1.2, 2.5) > 366^\circ.5$.

If there are at most 5 points of weight 1/2 in $1.2 \leq \rho \leq 5/3$, then there must be at least 10 points of weight 1/2 in $1 \leq \rho \leq 1/2$. By Lemma 3.11, this implies that there are not 12 points of weight 1/2 within $1 \leq \rho \leq 1.4$. Therefore there are at least 4 points of weight 1/2 in $1.4 \leq \rho \leq 5/3$. We now show that this is not possible.

Sublemma 3.15.2 *It is impossible to have 8 admissible points of weight 1 in $1.5 \leq \rho \leq 2.5$ and 4 admissible points of weight 1/2 in $1.35 \leq \rho \leq 5/3$.*

Proof of sublemma. No configuration of the 12 points exists such that the angles subtended sum to 360° or less. The smallest value possible occurs when the 4 points of weight 1/2 are consecutive. The sum of angles in this case is $7\Phi_{1.5}(1.5, 2.5) + 2\Phi_{1.5}(1.4, 2.5) + 3\Phi_{2/3}(1.4, 5/3) > 366^{\circ}.7$.

This proves the lemma and therefore case 3.

□

Lemma 3.16 (Case 4) *It is impossible to have 7 admissible points of weight 1 and 17 admissible points of weight 1/2.*

Sublemma 3.16.1 *It is impossible to have 7 admissible points of weight 1 in $1.5 \leq \rho \leq 2.5$ and 8 admissible points of weight 1/2 in $1.2 \leq \rho \leq 5/3$.*

Proof of sublemma. No configuration of the 15 points exists such that the angles subtended sum to 360° or less. The smallest value possible occurs when the 7 points of weight 1 are each between points of weight 1/2. The sum of angles in this case is $14\Phi_{1.5}(1.2, 2.5) + \Phi_{2/3}(1.2, 5/3) > 368^{\circ}.6$.

If there are at most 7 points of weight 1/2 in $1.2 \leq \rho \leq 5/3$, then there must be at least 10 points of weight 1/2 in $1 \leq \rho \leq 1/2$. By Lemma 3.11, this implies that there are not 12 points of weight 1/2 within $1 \leq \rho \leq 1.4$. Therefore there are at least 6 points of weight 1/2 in $1.4 \leq \rho \leq 5/3$. We now show that this is not possible.

Sublemma 3.16.2 *It is impossible to have 7 admissible points of weight 1 in* $1.5 \leq \rho \leq 2.5$ *and 6 admissible points of weight 1/2 in* $1.4 \leq \rho \leq 5/3$.

Proof of sublemma. No configuration of the 13 points exists such that the angles subtended sum to 360° or less. The smallest value possible occurs when the 6 points of weight 1/2 are consecutive. The sum of angles in this case is $6\Phi_{1.5}(1.5, 2.5) + 2\Phi_{1.5}(1.4, 2.5) + 5\Phi_{2/3}(1.4, 5/3) > 379°.2$.

This proves the lemma and therefore case 4.

□

Lemma 3.17 (Case 5) *It is impossible to have 6 admissible points of weight 1 and 19 admissible points of weight 1/2.*

Sublemma 3.17.1 *It is impossible to have 6 admissible points of weight 1 in* $1.5 \leq \rho \leq 2.5$ *and 10 admissible points of weight 1/2 in* $1.2 \leq \rho \leq 5/3$.

Proof of sublemma. No configuration of the 16 points exists such that the angles subtended sum to 360° or less. The smallest value possible occurs when the 6 points of weight 1 are each between points of weight 1/2. The sum of angles in this case is $12\Phi_{1.5}(1.2, 2.5) + 4\Phi_{2/3}(1.2, 5/3) > 376°.9$.

If there are at most 9 points of weight 1/2 in $1.2 \leq \rho \leq 5/3$, then there must be at least 10 points of weight 1/2 in $1 \leq \rho \leq 1/2$.

By Lemma 3.11, this implies that there are not 12 points of weight 1/2 within $1 \leq \rho \leq 1.4$. Therefore there are at least 8 points of weight 1/2 in $1.4 \leq \rho \leq 5/3$. We now show that this is not possible.

Sublemma 3.17.2 *It is impossible to have 6 admissible points of weight 1 in* $1.5 \leq \rho \leq 2.5$ *and 8 admissible points of weight 1/2 in* $1.4 \leq \rho \leq 5/3$.

Proof of sublemma. No configuration of the 14 points exists such that the angles subtended sum to 360° or less. The smallest value possible occurs when the 9 points of weight 1/2 are consecutive. The sum of angles in this case is $5\Phi_{1.5}(1.5, 2.5) + 2\Phi_{1.5}(1.4, 2.5) + 7\Phi_{2/3}(1.4, 5/3) > 391°.7$.

This proves the lemma and therefore case 5.

□

Lemma 3.18 (Case 6) *It is impossible to have 26 admissible points.*

With reasoning parallel to the proof of Lemma 3.7, we may assume that all 26 points are of weight 1/2, creating a set Δ^* of 26 admissible points that all lie in $1 \leq \rho \leq 5/3$. We make two important observations:

$$16\Phi_{2/3}(1.4, 5/3) > 368°.0$$
$$14\Phi_{2/3}(1, 1.4) > 364°.0$$

There are only three cases in which we can achieve 26 admissible points. Either there are 15 points in $1.4 \leq \rho \leq 5/3$ and 11 in $1 \leq \rho \leq 1.4$, or there are 14 and 12, or 13 and 13. We will show that none of these are possible.

Sublemma 3.18.1 *It is impossible to have 15 admissible points in $1.4 \leq \rho \leq 5/3$ and 11 in $1 \leq \rho \leq 1.4$.*

Proof of sublemma. Note that no points are in $1.3 \leq \rho \leq 1.4$, since $14\Phi_{2/3}(1.4, 5/3) + 2\Phi_{2/3}(1.3, 5/3) > 365°.6$. Therefore, by Lemma 3.10, 2 of the points in $1 \leq \rho \leq 1.4$ must lie in $1.15 \leq \rho \leq 1.3$. If the two points each lie between points in $1.4 \leq \rho \leq 5/3$, then the sum of the angles subtended is $13\Phi_{2/3}(1.4, 5/3) + 4\Phi_{2/3}(1.15, 5/3) > 369°.0$. If the two are consecutive, then the sum of the angles subtended is $14\Phi_{2/3}(1.4, 5/3) + 2\Phi_{2/3}(1.15, 5/3) + \Phi_{2/3}(1.15, 1.3) > 386°.7$.

Sublemma 3.18.2 *It is impossible to have 14 admissible points in $1.4 \leq \rho \leq 5/3$ and 12 in $1 \leq \rho \leq 1.4$.*

Proof of sublemma. By Lemma 3.11, 3 of the points in $1 \leq \rho \leq 1.4$ must lie in $1.2 \leq \rho \leq 1.4$. If the three points each lie between points in $1.4 \leq \rho \leq 5/3$, then the sum of the angles subtended is $11\Phi_{2/3}(1.4, 5/3) + 6\Phi_{2/3}(1.2, 5/3) > 368^\circ.8$. If two are consecutive and the third lies between points in $1.4 \leq \rho \leq 5/3$, then the sum is $12\Phi_{2/3}(1.4, 5/3) + 4\Phi_{2/3}(1.2, 5/3) + \Phi_{2/3}(1.2, 1.4) > 380^\circ.7$. If all three are consecutive, then the sum is $13\Phi_{2/3}(1.4, 5/3) + 2\Phi_{2/3}(1.2, 5/3) + 2\Phi_{2/3}(1.2, 1.4) > 392^\circ.6$.

Sublemma 3.18.3 *It is impossible to have 13 admissible points in $1.4 \leq \rho \leq 5/3$ and 13 in $1 \leq \rho \leq 1.4$.*

Proof of sublemma. By Lemma 3.11, 4 of the points in $1 \leq \rho \leq 1.4$ must lie in $1.2 \leq \rho \leq 1.4$. The minimum sum of the angles subtended occurs when these four points each lie between points in $1.4 \leq \rho \leq 5/3$. In this case, the sum of the angles is $11\Phi_{2/3}(1.4, 5/3) + 6\Phi_{2/3}(1.2, 5/3) > 368^\circ.8$, so this case is also impossible.

This proves the lemma and therefore case 6.

□

This completes the proof of Theorem 3.5, which implies that no node in the WSIG has outgoing edges whose weights sum to greater than 15. Since by Lemma 3.3 the total weight of the WSIG equals the number of edges in the SIG, we have proven Theorem 3.2, which states that no sphere of influence graph of n vertices in the Euclidean plane contains more than $15n$ edges.

Chapter 4

Conclusion

We have presented new bounds on the maximum size of the sphere of influence graph (SIG) both in d-dimensional infinite-order Minkowski space and in the Euclidean plane. In M_∞^d, there exist SIGs whose sizes asymptotically approach $(2^{2d+2} - 3d - 4)n/9$, and all SIGs contain fewer than $(2^{2d-1} - 2^{d-1})n$ edges. We have also shown that no SIG in the Euclidean plane contains more than $15n$ edges.

We leave several problems open. The first four concern M_∞^d-SIGs.

Open Problem 4.1 *Does there exist a function $c(d)$ less than $(2^{2d-1} - 2^{d-1})$ such that no M_∞^d-SIG has $c(d)n$ edges or greater?*

Open Problem 4.2 *Do there exist M_∞^d-SIGs of n vertices with more than* $(2^{2d+2} - 3d - 4)n/9$ *edges?*

Open Problem 4.3 *Does there exist a d-dimensional lattice whose M_∞^d-SIG contains more edges than T_d? If so, what is the optimal lattice?*

Given the difficulties in many of the covering and kissing problems[1] in higher dimensions, Open Problem 4.3 may be quite difficult.

[1]The reader is referred to [CS93] for an excellent book on these problems.

Open Problem 4.4 *We have discussed only briefly the* closed sphere of influence graph, *in which the spheres of influence are closed balls. Michael and Quint [MQ94a, MQ94b] present an upper bound of* $(5^d - (3/2))n$ *in* M_∞^d*; the best known lower bound is* $5^d n/2$*. Improve these bounds.*

The proof in Chapter 3 that no SIG in the Euclidean plane contains more than $15n$ edges is dependent on dividing the SIG edges into three discrete groups, each with its own weight. While it is conceivable that better results could result from dividing the edges into five or more groups, the resulting proof would be immense, ridden with several cases to analyze. Perhaps some continuous weighting scheme could be devised and the proof made more general to avoid the many cases. This leads us to another open problem.

Open Problem 4.5 *Does there exist a continuous weighting of the edges of the E-SIG which yields an upper bound better than 15n? Does it yield an elegant proof?*

Of course, the ultimate goal is to have the upper bound and the lower bound meet. We recall the conjecture made by David Avis on the maximum size of the sphere of influence graph in the Euclidean Plane.

Conjecture 4.6 (Avis) *No E-SIG or closed E-SIG contains more than 9n edges.*

We close this discussion on sphere of influence graphs with the beautifully simple yet elusive open problem that has inspired this research.

Open Problem 4.7 *Prove or disprove Conjecture 4.6.*

Appendix A

A review of lattices

This appendix provides a brief review of lattices and generating bases. Any introductory text on linear algebra should provide a more thorough coverage.

We will start with the definition of a *basis*. Given a vector space $\mathcal{X}$, a basis is a set $\mathcal{B}$ of vectors such that any point in $\mathcal{X}$ can be uniquely defined by a linear combination of vectors in $\mathcal{B}$. (We then call the vectors linearly independent.) For example, in the plane, the set $\{(1,0),(0,1)\}$ is the usual basis. However, we could also use the basis $\mathcal{B} = \{(0,1),(1,1)\}$. Here the point normally labelled $(6,8)$ can be defined uniquely as $-2(0,1)+8(1,1)$.

Note that in an m dimensional space, any basis has exactly m vectors. For example, in the plane the set $\{(0,1),(1,0),(1,1)\}$ can't be a basis since we can express $(6,8)$ as $-2(0,1)+8(1,1)$, or $-3(0,1)+7(1,0)+(1,1)$, or any other proper summation. (There are infinitely many possibilities.) Not just any m vectors will do; the set $\{(0,1),(0,2)\}$ will not allow us to express all vectors in the plane. The vectors $(0,1)$ and $(0,2)$ are not linearly independent.

A *lattice* is the set of all points that can be expressed as sums of *integral* multiples of vectors of a given basis $\mathcal{B}$. The basis $\mathcal{B}$ is said to be the *generating basis* of the lattice. Figure A.1 is an example.

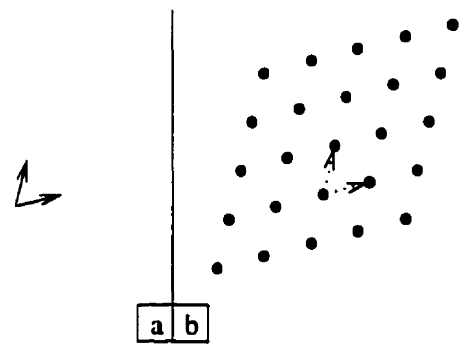

Figure A.1: (a) A generating basis and (b) its lattice.

Bibliography

[AH85] David Avis and Joe Horton. Remarks on the sphere-of-influence graph. *Discrete Geometry and Convexity*, 440:323–327, 1985.

[BE51] Paul Bateman and Paul Erdös. Geometrical extrema suggested by a lemma of Besicovitch. *American Mathematical Monthly*, 58:306–314, 1951.

[Bes45] Abram S. Besicovitch. A general form of the covering principle and relative differentiation of additive functions. *Proceedings of the Cambridge Philosophical Society*, 41:103–110, 1945.

[CS93] John H. Conway and N.J.A. Sloane. *Sphere Packings, Lattices, and Groups (Second Edition)*. Springer-Verlag, 1993.

[Dwy93] Rex A. Dwyer. The expected size of the sphere-of-influence graph. Technical Report TR-93-03, North Carolina State University, 1993.

[GPS94] Leonidas Guibas, János Pach, and Micha Sharir. Sphere-of-influence graphs in higher dimensions. *Colloquia Mathematica Societatis János Bolyai*, Intuitive Geometry, 63:131–137, 1994.

[JT92] Jerzy W. Jaromczyk and Godfried T. Toussaint. Relative neighborhood graphs and their relatives. *Proceedings of the IEEE*, 80(9):1502–1517, 1992.

[MQ94a] T. S. Michael and T. Quint. Sphere of influence graphs: a survey. *Congressus Numerantium*, 105:153–160, 1994.

[MQ94b] T. S. Michael and T. Quint. Sphere of influence graphs: edge density and clique size. *Mathematical and Computer Modelling*, 20:19–24, 1994.

[Rei48] E. R. Reifenberg. A problem on circles. *Mathematical Gazette*, 32:290–292, 1948.

[Tou80] Godfried T. Toussaint. Pattern recognition and geometric complexity. In *Proceedings of the 5th International Conference on Pattern Recognition*, pages 1324–1347, Miami Beach, 1980.

[Tou81] Godfried T. Toussaint. Computational geometric problems in pattern recognition. *Pattern Recognition Theory and Applications*, pages 73–91, 1981.

[Tou88] Godfried T. Toussaint. A graph-theoretic primal sketch. *Computational Morphology*, pages 229–260, 1988.