A space-filling complete graph

J. Ginsburg and V. Linek
University of Winnipeg
Winnipeg, Manitoba
Canada R3B 2E9

August 4, 1999

Abstract

We show that there is a straight line embedding of the complete graph $K_{\mathbf{c}}$ into $\mathbf{R}^3$ which is space-filling: every point of $\mathbf{R}^3$ is either one of the vertices of $K_{\mathbf{c}}$, or lies on exactly one straight line segment joining two of the vertices.

1 Introduction and preliminaries

It is a familiar and basic fact to any student of graph theory that the complete graph on 5 vertices is not planar. That is, it is impossible to choose 5 points in the plane and join every pair of them by a straight line segment in such a way that no two of the line segments cross. Denoting the complete graph on m vertices by K_m, we also say that K_5 cannot be *embedded* in the Euclidean plane $\mathbf{R}^2$. One can of course also consider the drawing and embedding of graphs in 3-dimensional Euclidean space. Here one can show that, for any positive integer m, it is possible to draw or represent K_m in $\mathbf{R}^3$ in such a way that the edges of K_m are represented by straight line segments, no two of which intersect, except at their end-points when appropriate. In fact, this can even be done for the complete graph on a continuum of vertices. That is, let $\mathbf{c}$ denote the cardinal number of the set of real numbers $\mathbf{R}$. It is possible to choose a set of $\mathbf{c}$ points in $\mathbf{R}^3$ and draw straight line segments in $\mathbf{R}^3$ connecting every pair of these points in such a way that no two of these straight line segments intersect (except possibly at their end-points). In this note we will show that such a drawing of the complete graph $K_{\mathbf{c}}$ can be constucted so as to *fill* all of $\mathbf{R}^3$; that is, in such a way that every point of $\mathbf{R}^3$ is either one of the vertices of $K_{\mathbf{c}}$, or lies on exactly one straight line segment joining two of the vertices. While this might be thought of as a kind of discrete counterpart to the notion of a space-filling curve, it seems to us to be more akin to certain kinds of tiling and covering problems, as studied for example in [3] and [2].

Let us now introduce the terminology and notation to be used. We refer the reader to [6] and [1] respectively for any additional background material on set theory and graph theory respectively. As usual we denote the cardinality of a set S by $|S|$. Infinite cardinal numbers are thought of as initial ordinals. In this paper, a ***graph*** G is a pair (V, E), where V is a set and where E is a set of two-element subsets of V. As usual, the elements of V are called the *vertices* of G and the elements of E are called the *edges* of G. If E is the set of *all* two-element subsets of V then G is called the *complete graph* on the vertex set V. If $|V| = m$ we denote the complete graph on the vertex set V by K_m. If $G = (V, E)$ and if S is a subset of V, we say that S is a *covering* for G if every edge of G intersects S. A *matching* in G is a pairwise disjoint set of edges of G. If W is a subset of V then the *subgraph of G induced by W* is the graph $H = (W, F)$ where $F = \{ \{u, v\} \in E \mid u \in W \text{ and } v \in W\}$.

If n is a positive integer and if p and q are two distinct points in $\mathbf{R}^n$, we denote the *straight line* in $\mathbf{R}^n$ passing through p and q by $L_{p,q}$. The closed segment of this line extending from p to q is denoted by $[p, q]$. We will denote the open line segment joining p and q by (p, q). That is, $(p, q) = [p, q] - \{p, q\}$. If S is any subset of $\mathbf{R}^n$ we let $\langle S \rangle$ denote the *affine hull* of S in $\mathbf{R}^n$; this is the smallest subset T of $\mathbf{R}^n$ which contains S and which contains the straight line joining any two of its points. For $k < n$, a translate of a k-dimensional vector subspace of $\mathbf{R}^n$ is called a *k-flat* in $\mathbf{R}^n$.

Let $G = (V, E)$ be a graph and let n be a positive integer. A *straight line embedding* of G in $\mathbf{R}^n$ is a one-to-one function $f : V \to \mathbf{R}^n$ such that $(f(u_1), f(v_1)) \cap (f(u_2), f(v_2)) = \phi$ for any two distinct edges $\{u_1, v_1\}$ and $\{u_2, v_2\}$ of G. In other words, a straight line embedding of G in $\mathbf{R}^n$ is a "drawing" of the graph G in which the vertices of G are represented by points in $\mathbf{R}^n$, and each edge $\{u, v\}$ of G is represented by the straight line segment $[f(u), f(v)]$ joining the points corresponding to u and v. The line segments representing the edges do not intersect except possibly at their end-points (when two edges are incident with a common vertex). If f is a straight line embedding of G in $\mathbf{R}^n$ and if $\mathbf{R}^n = f(V) \cup \left(\bigcup_{\{u,v\}\in E}[f(u), f(v)]\right)$ we say that f is a *space-filling embedding* of G in $\mathbf{R}^n$.

We remark that straight line embeddings of infinite graphs in $\mathbf{R}^2$ (that is, planar representations of infinite graphs), and infinite graphs which admit such embeddings, have been studied and characterized in [8], to which we refer the interested reader.

2 The construction of a space-filling straight line embedding

We will now proceed to show that there is a space-filling straight line embedding of K_c in $\mathbf{R}^3$. The following simple lemma is a key ingredient in our construction.

Lemma 2.1 Let n be a positive integer. Then $\mathbf{R}^n$ is not the union of $< c$ k-flats each of dimension less than n.

Proof: Let $\mathcal{D}$ be a decomposition of $\mathbf{R}^n$ into c parallel copies of $(n-1)$-flats. Suppose now that $\mathbf{R}^n = \bigcup_{a\in A} F_a$ where $|A| < c$ and each F_a is a k-flat for some $k \leq n-1$. Then since $|A| < c$, there is some $(n-1)$-flat K belonging to $\mathcal{D}$ which is different from all of the F_a. If $F'_a = F_a \cap K$ then F'_a is either empty or is a k-flat of dimension at most $n-2$. Since $K = \bigcup_{a\in A} F'_a$, this shows that $\mathbf{R}^{n-1}$ is the union of $< c$ k-flats each of dimension less than $n-1$. The lemma now follows by induction. □

Theorem 2.2 Let $G = (V, E)$ be a graph with $|V| = c$ such that G contains a matching of size c. Then there is a space-filling straight line embedding of G in $\mathbf{R}^3$.

Proof: Let the points of $\mathbf{R}^3$ be well-ordered as $\{p_\alpha | \alpha < c\}$, and the vertices of G as $\{v_\alpha | \alpha < c\}$. Our embedding of G will be constructed by transfinite induction. We will construct, for every $\alpha < c$, an induced subgraph $H_\alpha = (W_\alpha, F_\alpha)$ of G, and a straight line embedding f_α of H_α in $\mathbf{R}^3$ such that for all $\alpha < c$:
(i) $|W_\alpha| \leq |\alpha| + \omega$.
(ii) $v_\alpha \in W_\alpha$.
(iii) If $\beta < \alpha$ then $W_\beta \subseteq W_\alpha$ and $f_\beta = f_\alpha | W_\beta$.
(iv) $p_\alpha \in f_\alpha(W_\alpha) \cup \left(\bigcup_{\{u,v\}\in F_\alpha} [f_\alpha(u), f_\alpha(v)]\right)$.

To begin the induction we let H_0 be the subgraph of G induced by the one-element set $W_0 = \{v_0\}$ and we define $f_0(v_0) = p_0$.

Now suppose that $\gamma < c$ and that for every $\alpha < \gamma$ we have constructed an induced subgraph H_α of G and a straight line embedding f_α of H_α into $\mathbf{R}^3$ satisfying the above properties for all $\alpha < \gamma$. We now describe how to construct H_γ and f_γ. We will find it convenient to use the following terminology : If f is a straight line embedding of a graph (W, F) in $\mathbf{R}^3$ and if p is a point of $\mathbf{R}^3$, we will say that p *is covered by the embedding* f if $p \in f(W) \cup \left(\bigcup_{\{u,v\}\in F} [f(u), f(v)]\right)$.

First we let $U_\gamma = \bigcup_{\alpha<\gamma} W_\alpha$. We will eventually define W_γ to be a certain subset of V containing U_γ. Our induction hypothesis implies that $g_\gamma = \bigcup_{\alpha<\gamma} f_\alpha$ is a straight line embedding of the subgraph of G induced by U_γ.

If the point p_γ is covered by any of the embeddings f_α for some $\alpha < \gamma$ (that is, covered by g_γ), we can omit the following argument and proceed directly to the second part of the proof below (taking V_γ to be U_γ and h_γ to be g_γ). So suppose that the point p_γ is *not* covered by any of the embeddings f_α for $\alpha < \gamma$.

Let $\mathcal{F}$ be the collection of all k-flats $\langle g_\gamma(x), g_\gamma(y), g_\gamma(z)\rangle$ and $\langle g_\gamma(x), g_\gamma(y), p_\gamma\rangle$ where x, y, and z are vertices of U_γ. These flats include all planes determined by three non-collinear points of $g_\gamma(U_\gamma)$, all lines through two or three points of $g_\gamma(U_\gamma)$, and all of the points of $g_\gamma(U_\gamma)$, as well as all the lines and planes through p_γ and one or two points of $g_\gamma(U_\gamma)$.

Now condition (i) clearly implies that $|U_\gamma| \leq |\gamma| + \omega < \mathfrak{c}$. Therefore there are less than $\mathfrak{c}$ flats belonging to the collection $\mathcal{F}$. By the preceding lemma, these flats cannot cover all of $\mathbf{R}^3$ and so there is a point q in $\mathbf{R}^3$ which is not in any of them. Let L be the line in $\mathbf{R}^3$ passing through p_γ and q. Then L is not contained in any of the flats belonging to $\mathcal{F}$. Therefore, for every flat F belonging to $\mathcal{F}$, we have that $L \cap F$ is either empty or consists of a single point. Since $|\mathcal{F}| < \mathfrak{c}$, this implies that $|L \cap (\bigcup_{F\in\mathcal{F}} F)| < \mathfrak{c}$. Since the line L contains $\mathfrak{c}$ points on either side of the point p_γ, there are points p_σ and p_τ in $L - (\bigcup_{F\in\mathcal{F}} F)$ such that $p_\gamma \in (p_\sigma, p_\tau)$.

Since the graph G has contains a matching of size $\mathfrak{c}$ there is an edge $\{v_\kappa, v_\lambda\}$ in G which is disjoint from the set U_γ. Let $V_\gamma = U_\gamma \cup \{v_\kappa, v_\lambda\}$ and let h_γ be the function from V_γ to $\mathbf{R}^3$ which extends the function g_γ on U_γ and such that $h_\gamma(v_\kappa) = p_\sigma$ and $h_\gamma(v_\lambda) = p_\tau$.

We now show that h_γ is a straight line embedding into $\mathbf{R}^3$ of the subgraph of G induced by V_γ. Consider any two distinct edges e_1 and e_2 of this subgraph, at least one of which is incident with one of the vertices in the set $\{v_\kappa, v_\lambda\}$. We consider all the various possibilities.

Case 1: If these edges are of the form $e_1 = \{v_\kappa, u\}$ and $e_2 = \{v_\kappa, v\}$ where u and v are vertices of U_γ then $(h_\gamma(v_\kappa), h_\gamma(u))$ and $(h_\gamma(v_\kappa), h_\gamma(v))$ must be disjoint: otherwise the point $p_\sigma = h_\gamma(v_\kappa)$ would be on the line $\langle g_\gamma(u), g_\gamma(v)\rangle$, contrary to the way p_σ was chosen. An identical argument applies to any two edges of the form $e_1 = \{v_\lambda, u\}$ and $e_2 = \{v_\lambda, v\}$ where u and v are vertices of U_γ .

Case 2: Suppose that $e_1 = \{v_\kappa, u\}$ and $e_2 = \{v, w\}$ where u, v, w are vertices of U_γ. In this case (see Figure 1a), if $(h_\gamma(v_\kappa), h_\gamma(u)) \cap (h_\gamma(v), h_\gamma(w)) \neq \phi$ then $p_\sigma = h_\gamma(v_\kappa)$ must be in $\langle g_\gamma(u), g_\gamma(v), g_\gamma(w)\rangle$, again contrary to the choice of p_σ .

Case 3: Suppose that $e_1 = \{v_\kappa, u\}$ and $e_2 = \{v_\lambda, v\}$ where u, v are vertices of U_γ. Suppose that $(h_\gamma(v_\kappa), h_\gamma(u)) \cap (h_\gamma(v_\lambda), h_\gamma(v)) \neq \phi$. (see Figure 1b) Since neither of p_σ or p_τ is in the flat $\langle g_\gamma(u), p_\gamma\rangle$ we must have $u \neq v$. Since the point p_γ was not covered by any of the embeddings f_α for $\alpha < \gamma$, we know that p_γ is not on the closed line segment joining $g_\gamma(u)$ and $g_\gamma(v)$ and therefore the points $p_\gamma, g_\gamma(u), g_\gamma(v)$ cannot be collinear. Since $(p_\sigma, g_\gamma(u)) \cap (p_\tau, g_\gamma(v)) \neq \phi$, it follows that both p_σ and p_τ are in the plane $\langle g_\gamma(u), g_\gamma(v), p_\gamma\rangle$, a contradiction.

Case 4: Suppose that $e_1 = \{v_\kappa, v_\lambda\}$ and $e_2 = \{u, v\}$ where u, v are vertices of

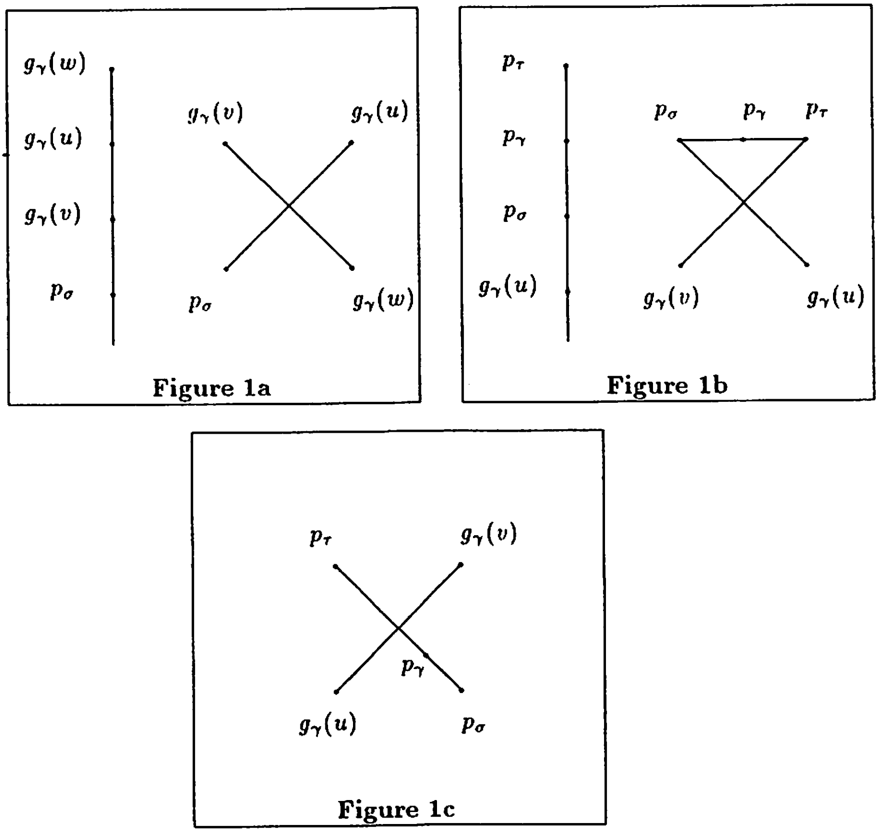

Figure 1: The various cases.

U_γ. In this case (see Figure 1c), if $(h_\gamma(v_\kappa), h_\gamma(v_\lambda)) \cap (h_\gamma(u), h_\gamma(v)) \neq \phi$ then this intersection cannot be at the point p_γ since p_γ is not covered by any of the embeddings f_α for $\alpha < \gamma$. Therefore this would again imply that both p_σ and p_τ are in the plane $\langle g_\gamma(u), g_\gamma(v), p_\gamma \rangle$, a contradiction.

Case 5: Suppose that $e_1 = \{v_\kappa, v_\lambda\}$ and $e_2 = \{v_\kappa, v\}$ where v is a vertex of U_γ. In this case, we see that the intersection $(h_\gamma(v_\kappa), h_\gamma(v_\lambda)) \cap (h_\gamma(v_\kappa), h_\gamma(v))$ must be empty since, otherwise, the points p_σ, p_τ and $g_\gamma(v)$ would be collinear. Since p_γ is also on the line through p_σ and p_τ, this implies that both p_σ, p_τ are in the flat $\langle g_\gamma(v), p_\gamma \rangle$, a contradiction. An identical argument applies when κ and λ are interchanged.

Thus we have that h_γ is a straight line embedding into $\mathbf{R}^3$ of the subgraph of G induced by V_γ. Furthermore, h_γ extends f_α for all $\alpha < \gamma$ and p_γ is covered by the embedding h_γ.

Now, if the vertex v_γ belongs to the set V_γ we let $W_\gamma = V_\gamma$ and we let $f_\gamma = h_\gamma$ to complete the inductive step. If not, we consider the collection $\mathcal{F}'$ of all the flats $\langle h_\gamma(x), h_\gamma(y), h_\gamma(z)\rangle$ and where x, y, and z are vertices of V_γ. There are less than $\mathfrak{c}$ flats belonging to this collection and so, by the lemma, there is a point p_ξ in $\mathbf{R}^3$ which is not in any of them. We let $W_\gamma = V_\gamma \cup \{v_\gamma\}$ and we let f_γ be the function defined on W_γ which extends h_γ and for which $f_\gamma(v_\gamma) = p_\xi$. The same reasoning we applied in considering the various cases above shows that f_γ is a straight line embedding into $\mathbf{R}^3$ of the subgraph of G induced by W_γ. It is clear that conditions (i), (ii), (iii) and (iv) above are satisfied. This completes the inductive step.

Now, let $f = \bigcup_{\alpha<\mathfrak{c}} f_\alpha$. Conditions (i), (ii), (iii) and (iv) imply that f is a space-filling straight line embedding of G in $\mathbf{R}^3$. □

Since the graph $K_\mathfrak{c}$ obviously has a matching of cardinality $\mathfrak{c}$, we obtain the following as an immediate corollary.

Corollary 2.3 There is a space-filling straight line embedding of $K_\mathfrak{c}$ in $\mathbf{R}^3$.

One naturally wonders if there is a space filling embedding of $K_\mathfrak{c}$ that does not require the Axiom of Choice. For instance it is known (see [4]) that $K_\mathfrak{c}$ can be embedded in 3-space by letting the vertices be the points on the moment curve (t, t^2, t^3), $t \in \mathbf{R}$. For a set $S \subset \mathbf{R}^3$ let the *envelope* of S be the union of all closed line segments joining pairs of points of S. Any point on (p, q), where p, q are distinct points on the curve, is an interior point of the envelope of the moment curve, as may be seen by applying the Inverse Function Theorem. Intuitively we see this by constructing tangent line segments at p and q, which are skew, hence their envelope is a tetrahedron with non-empty interior. The envelope of two short arcs constructed at p and q will have a similar shape, and so also have non-empty interior. The envelope of the moment curve does not fill all of space, but it is an interesting question if a piecewise smooth curve can give a space filling embedding of $K_\mathfrak{c}$.

Our use of the assumption concerning the cardinality of a matching in a graph and some related ideas will be discussed in the next section.

3 On the cardinality of covering sets and a related question

The space-filling embedding in 2.3 gives a covering by a special family of non-crossing closed line segments (for which, if $[a, b]$ and $[c, d]$ belong to the family of line segments and $a \neq c$ then so does $[a, c]$). This kind of covering should be contrasted with the work in [3] and [2] where the authors study the possibility of covering spheres and planes by *congruent* arcs. In the case of 2.3, we note that we cannot possibly have all of the line segments congruent. In fact, the lengths of the line segments in any space-filling embedding of $K_{\mathfrak{c}}$ cannot even be bounded above: If all their lengths were $\leq r$, choose any point p of $\mathbf{R}^3$ which corresponds to a vertex of $K_{\mathfrak{c}}$. Since any other point q of $\mathbf{R}^3$ which corresponds to a vertex of $K_{\mathfrak{c}}$ is joined by one of the line segments to p, this implies that all such points q are no more than r units away from p. Since any other point of $\mathbf{R}^3$ lies on a line segment between two such points q_1 and q_2, this would imply that *all* points of $\mathbf{R}^3$ are within r units of p. It is also of interest to note that neither can the lengths of the line segments in any space-filling embedding of $K_{\mathfrak{c}}$ be bounded below. For, let P denote the set of $\mathfrak{c}$ points in space corresponding to the vertices of $K_{\mathfrak{c}}$. Since $\mathbf{R}^3$ is the union of a countable number of compact sets, there must be (by the pigeon-hole principle) an uncountable number of points of P contained in a compact subset K of $\mathbf{R}^3$. This infinite set of points has a cluster point in K. This implies that there are points of P which are arbitrarily close together, and so the lengths of the line segments joining them are arbitrarily small.

In order to construct our space-filling embedding we made essential use of the assumption that our graph G contained a matching of size $\mathfrak{c}$. How essential is this assumption? Well, for one thing, it is not necessary for a graph G to satisfy this condition just because there is a space-filling straight line embedding of G in $\mathbf{R}^3$. Such a graph G with no matching of size $\mathfrak{c}$ necessarily has a covering of size $< \mathfrak{c}$. Indeed there is a graph G of cardinality $\mathfrak{c}$ which has a *countable* covering and for which there is a space-filling straight line embedding of G in $\mathbf{R}^3$. We thank Bill Sands for suggesting the following example, which is simpler than our original one. We partition $\mathbf{R}^3$ into a countable number of unit cubes $\{C_n \mid n \in \mathbf{N}\}$ in the usual way using a three-dimensional grid. These cubes intersect only on their boundaries. Let p_n be the center of the cube C_n. Let G be the graph which has a vertex v_n corresponding to each of the centers p_n, and a vertex for each point on the boundary of any of the cubes C_n. For the edges of G, we take an edge from each v_n to all the vertices which correspond to the points on the boundary of the cube C_n. Clearly G has a countable covering and there is a natural space-filling straight line embedding of G in $\mathbf{R}^3$. We note that an identical construction can be applied in the plane $\mathbf{R}^2$ (using unit squares) to give a graph G which has a countable covering and for which there is a straight line embedding of G in $\mathbf{R}^2$ which fills the entire plane.

Can one characterize those graphs G for which there is a space-filling straight line embedding of G in $\mathbf{R}^3$? The preceding example shows that the smallest cardinality of a covering in G does not by itself lead to such a characterization. We have not been able to find such a characterization.

Having seen an example of a graph G which has a countable covering and for which there is a space-filling straight line embedding of G in $\mathbf{R}^3$, we may wish to go one step further and ask whether there is a graph G which has a *finite* covering and for which there is a space-filling straight line embedding of G in $\mathbf{R}^3$. At first glance, it may seem intuitively clear that such a graph cannot exist. However, we have only been been able to establish a few very special cases of this problem : the case when the graph G has a covering of size at most 3, and the case when G has a special kind of covering K of size 4 (for which any two edges incident with two different vertices of K are disjoint). The general finite case has so far eluded us. We will return to discuss these special cases in $\mathbf{R}^3$ at the end of this section. For the moment, let us instead consider the corresponding question for filling the *plane*. Here we can show, as expected, that there is no such graph with a finite covering. Our argument here has a much different flavour than the discussion in Section 2 above and, as we will see, leads to an interesting related question concerning the plane.

We will make use of the following terminology and notation in the plane. Let A be a point in $\mathbf{R}^2$, and let θ be a number such that $0 \le \theta \le 2\pi$. We let $L_A(\theta)$ denote the unbounded straight line ray in $\mathbf{R}^2$ extending from A in a direction which makes an angle of θ radians with the horizontal. If r is any non-negative real number, we let $L_A(\theta, r)$ denote the set of points on the ray $L_A(\theta)$ whose distance from A is at most r. Note that if A has coordinates (a_1, a_2), this is the same as the closed line segment $[A, B]$ where $B = (a_1 + r\cos\theta, a_2 + r\sin\theta)$. When $r = 0$, $L_A(\theta, r)$ just consists of the single point A. If θ_1 and θ_2 are two numbers such that $0 \le \theta_1 < \theta_2 \le 2\pi$, we let $C_A(\theta_1, \theta_2) = \bigcup_{\theta_1 \le \theta \le \theta_2} L_A(\theta)$. The set $C_A(\theta_1, \theta_2)$ is called a *sector* at A. The set of points in the sector $C_A(\theta_1, \theta_2)$ whose distance from A is $\ge r$ is denoted by $C_A^r(\theta_1, \theta_2)$ and is referred to as a *tail* of the sector $C_A(\theta_1, \theta_2)$.

It is easy to see that, if A and B are any two points in the plane, then any tail of a sector at A contains a tail of a sector at B (see Figure 2).

Let A be a point in the plane $\mathbf{R}^2$ and let g be a non-negative function defined on $[0, 2\pi]$ with $g(0) = g(2\pi)$. The set $\bigcup_{0 \le \theta \le 2\pi} L_A(\theta, g(\theta))$ is called

a *star at A* (corresponding to the function g) in $\mathbf{R}^2$ and will be denoted by $S_A(g)$. Thus a star at A contains a line segment in every direction from A: these segments can have all different lengths or even have length 0. A subset S of $\mathbf{R}^2$ is called a *star* if there is a point A and a function g such that $S \doteq S_A(g)$. These sets have also been called *radial sets* by other authors (for example, see

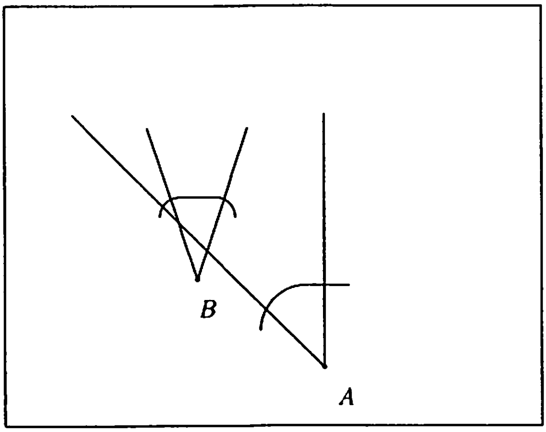

Figure 2: Tail of one sector contained in another.

[9].

Stars arise in an obvious way in connection with embeddings of graphs in $\mathbf{R}^2$. Let $G = (V, E)$ be a graph and let f be a straight line embedding of G in $\mathbf{R}^2$. Let u be a vertex of G. We will let $\mathcal{L}_u$ denote the set of all the line segments which correspond to the edges of G incident with u. That is, $\mathcal{L}_u = \{[f(u), f(v)] \mid \{u, v\} \in E\}$. For any two distinct vertices v_1 and v_2 of G which are adjacent to u in G, the two line segments $[f(u), f(v_1)]$ and $[f(u), f(v_2)]$ intersect only in the point $f(u)$ and so extend in different directions from this point. Thus the set $S_u = \{f(u)\} \cup \left(\bigcup_{\{u,v\} \in E}[f(u), f(v)]\right)$ is a star in $\mathbf{R}^2$ at the point $f(u)$.

It is clear that, if f is a space-filling straight line embedding of G in $\mathbf{R}^2$, and if U is a covering in G, then $\mathbf{R}^2 = \bigcup_{u \in U} S_u$.

Theorem 3.1 Let $G = (V, E)$ be a graph and suppose that G has a finite covering. Then there is no space-filling straight line embedding of G in $\mathbf{R}^2$.

Proof: We argue by contradiction. Suppose there is a space-filling straight line embedding f of G in $\mathbf{R}^2$. Let $U = \{u_1, u_2, ..., u_n\}$ be a finite covering of G, where n is some positive integer. As above, we consider the sets of lines $\mathcal{L}_{u_1}, \mathcal{L}_{u_2}, ..., \mathcal{L}_{u_n}$ which correspond to the set of edges incident with each of the vertices $u_1, u_2, ..., u_n$, and the stars $S_{u_1}, S_{u_2}, ..., S_{u_n}$ in $\mathbf{R}^2$ corresponding to these vertices. The union of these stars is all of $\mathbf{R}^2$. Let $g_1, g_2, ..., g_n$ be non-negative real-valued functions on $[0, 2\pi]$ such that $S_{u_i} = S_{f(u_i)}(g_i)$ for all $i = 1, 2, ..., n$.

We consider two possibilities. First, suppose that, for all $i = 1, 2, ..., n$

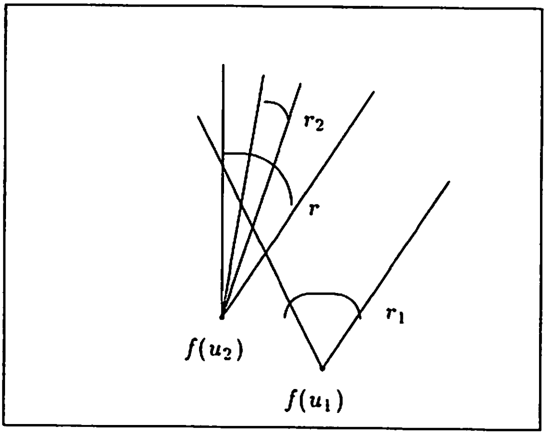

Figure 3:

and for every non-empty interval I contained in $[0, 2\pi]$, there is a non-empty interval $J \subseteq I$ such that g_i is bounded on J. Then there is a non-empty interval $J_1 = (\kappa_1, \lambda_1)$ and a positive real number r_1 such that $g_1(\theta) < r_1$ for all θ in J_1. The tail $C^{r_1}_{f(u_1)}(\kappa_1, \lambda_1)$ of the sector at $f(u_1)$ corresponding to the interval J_1 is disjoint from the star $S_{f(u_1)}(g_1)$. Let $T_1 = C^{r_1}_{f(u_1)}(\kappa_1, \lambda_1)$. Now there is a tail of a sector at $f(u_2)$ which is contained in the tail T_1. That is, there is an interval (κ, λ) and a positive r such that $C^{r}_{f(u_2)}(\kappa, \lambda) \subseteq T_1$. Because of our assumption, there is an interval $J_2 = (\kappa_2, \lambda_2)$ contained in the interval (κ, λ) and a positive number $r_2 > r$ such that $g_2(\theta) < r_2$ for all θ in J_1. Note that the tail $C^{r_2}_{f(u_2)}(\kappa_2, \lambda_2)$ is disjoint from the star $S_{f(u_2)}(g_2)$ and also from the star $S_{f(u_1)}(g_1)$, since $C^{r_2}_{f(u_2)}(\kappa_2, \lambda_2) \subseteq C^{r}_{f(u_2)}(\kappa, \lambda) \subseteq T_1$. Let $T_2 = C^{r_2}_{f(u_2)}(\kappa_2, \lambda_2)$. (see Figure 3) Continuing in this way, we can construct T_i for all $i = 1, 2, ..., n$ such that, for all i, T_i is a tail of a sector at the point $f(u_i)$ and T_i is disjoint from all of the stars $S_{f(u_1)}(g_1), S_{f(u_2)}(g_2), ..., S_{f(u_i)}(g_i)$. When we reach $i = n$, the tail T_n is disjoint from all of the stars
$S_{f(u_1)}(g_1), S_{f(u_2)}(g_2), ..., S_{f(u_n)}(g_n)$, contradicting the fact that these n stars cover the whole plane.

Thus we can assume that there is one of the functions g_i, say g_1, and a non-empty interval $I = (\kappa, \lambda)$, such that g_1 is unbounded on every non-empty interval J contained in I. Now, choose a non-empty interval $I_1 = (\kappa_1, \lambda_1)$ contained in I and so small that the sector $C_{f(u_1)}(\kappa_1, \lambda_1)$ contains none of the points $f(u_2), f(u_3), f(u_4), ..., f(u_n)$. Let R be a positive number such that $R > d(f(u_1), f(u_i))$ for all $i = 2, 3, ..., n$.(d denotes the usual Euclidean distance). Let θ be any number in the interval I_1 such that the ray $L_{f(u_1)}(\theta)$ is not parallel to any of the lines $L_{f(u_1), f(u_i)}$ for $i = 2, 3, ..., n$. Since g_1 is unbounded

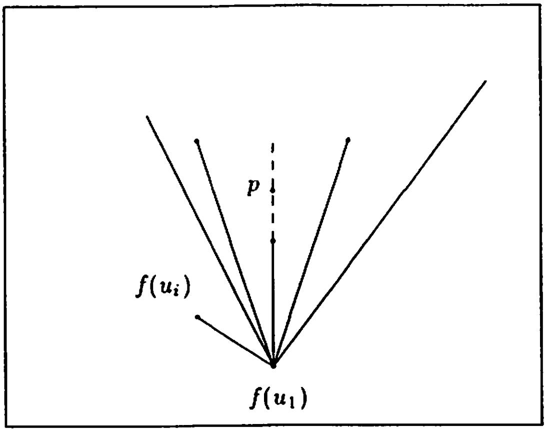

Figure 4:

on every non-empty interval contained in I there is a number θ_1 in the interval (κ_1, θ) and a number θ_2 in the interval (θ, λ_1) such that both $g_1(\theta_1)$ and $g_1(\theta_2)$ are larger than both of the numbers R and $g_1(\theta)$. Now, let p be any point which lies on the ray $L_{f(u_1)}(\theta)$ and whose distance from $f(u_1)$ is larger than $g_1(\theta)$ but smaller than both $g_1(\theta_1)$ and $g_1(\theta_2)$. Then p is not in the star $S_{f(u_1)}(g_1)$. Since the stars $S_{f(u_1)}(g_1), S_{f(u_2)}(g_2), \ldots, S_{f(u_n)}(g_n)$ cover the plane, there must be some $i \neq 1$ such that p is in the star corresponding to u_i. But (see Figure 4) any line segment extending from $f(u_i)$ which contains p obviously must meet one of the line segments $L_{f(u_1)}(\theta_1, g_1(\theta_1))$ or $L_{f(u_1)}(\theta_2, g_1(\theta_2))$ at a point in the interior of that line segment. This contradicts the fact that no line segment in the set $\mathcal{L}_{u_i}$ crosses any line segment in the set $\mathcal{L}_{u_1}$, since f is a straight line embedding. □

There is an interesting related question which arises very naturally in light of 3.1. Suppose that f is a straight line embedding of a graph G in $\mathbf{R}^2$. For any two vertices u and v of G, the line segments which make up the corresponding stars S_u and S_v do not intersect, except possibly at their end-points. What we have shown in the proof of 3.1 is that no finite number of such stars can cover the whole plane. Now, let us ignore graphs for the moment, and forget about the condition involving the non-intersection of the line segments making up the stars. Let us just think of stars in $\mathbf{R}^2$ in their own right. Can there be a finite number of stars in $\mathbf{R}^2$ whose union is all of $\mathbf{R}^2$? One's intuition suggests no. It is easy to see that two stars cannot cover $\mathbf{R}^2$. Also, a finite number of measurable stars cannot cover $\mathbf{R}^2$. This follows from the fact that the measure inside a big disk of a measurable star is small when compared to the measure of the disk. However, the general case has so far eluded us.

As mentioned above, we find ourselves in a much more unsatisfying state for $\mathbf{R}^3$. In this case we have only been able to establish the following very specialized result.

Theorem 3.2 Let $G = (V, E)$ be a graph. If G has a covering consisting of three or fewer vertices, or if G has a covering K consisting of four vertices such that any two edges of G incident with two different vertices of K are disjoint, then there is no space-filling straight line embedding of G in $\mathbf{R}^3$.

Proof: We will sketch the argument. Suppose there is a space-filling straight line embedding f of G in $\mathbf{R}^3$. If G has a cover K with $|K| \leq 3$ then the result follows by applying 3.1 to a plane containing the points of $f(K)$. In the second case, we can assume that the four vertices of $f(K)$ are not coplanar, otherwise again we can apply 3.1 to a plane containing $f(K)$. So, letting $K = \{u_1, u_2, u_3, u_4\}$, and letting $p_i = f(u_i)$, we can assume that p_1, p_2, p_3, p_4 are the vertices of a tetrahedron T in $\mathbf{R}^3$. Our assumption concerning K implies that, for every pair i, j with $i \neq j$ there is a point on the line segment (p_i, p_j) which is covered by a line $\ell_{i,j}$ through one of the other points p_k. The intersection $\ell^T_{i,j}$ of this line with T is contained in one of the faces of T. Since there are six such line segments (p_i, p_j) and T has only four faces, there must be two of the line segments $\ell^T_{i,j}$ which are contained in the same face of T and which cross. □

4 Addendum

Recently P. Komjáth [7] has shown that three stars can cover the plane if the Continuum Hypothesis holds. In fact the result is shown for certain subsets of stars called *clouds*, a cloud around a being a set of points that intersects every line through a in a finite set, and this result is shown to be equivalent to the Continuum Hypothesis.

5 Acknowledgement

The authors thank the anonymous referee for many helpful comments and suggestions, and for bringing [5] to their attention.

References

[1] J. A. Bondy and U. S. R. Murty. *Graph Theory with Applications*, Elsevier North Holland, Inc., 1976.

[2] H. G. Bothe. Parkettierungen der Ebene durch Kongruente Bogen. Ph.D. Thesis, Berlin, 1964.

[3] J. H. Conway and H. T. Croft. Covering a sphere with congruent great-circle arcs, *Proc. Camb. Phil. Soc.* 60 (1964): 787-800.

[4] B. Grünbaum. *Convex Polytopes*, John Wiley and Sons, 1967.

[5] B. Grünbaum and G. C. Shephard. *Tilings and Patterns*, W. H. Freeman and Company, 1986.

[6] K. Hrbacek and T. Jech. *Introduction to Set Theory*, Marcel Dekker, Inc., New York, 1984.

[7] P. Komjáth. Three clouds may cover the plane. Preprint, 3 pages.

[8] C. Thomassen. Straight line representations of infinite planar graphs. *J. London Math. Soc.* 16, No. 2 (1977): 411-423.

[9] S. Willard. *General Topology*, Addison-Wesley Publishing Co., 1970.