A Bounded Compactness Theorem for L^1-Embeddability of Metric Spaces in the Plane

Seth M. Malitz and Jerome I. Malitz
Computer Science Department
University of Massachusetts

COINS Technical Report 92-25
February 26, 1992

A Bounded Compactness Theorem for L^1-Embeddability of Metric Spaces in the Plane

Seth M. Malitz*
Dept. of Computer Science
Univ. of Massachusetts
Amherst, MA 01003
malitz@cs.umass.edu

Jerome I. Malitz
Dept. of Mathematics
Univ. of Colorado
Boulder, CO 80309
malitz@boulder.colorado.edu

February 26, 1992

*Research partially supported by NSF Grant CCR-8909953

Abstract

Let M be a metric space. An L^1-*embedding* of M into Cartesian k-space R^k is a distance-preserving map of M into R^k with the L^1-metric. Let $c(k)$ be the smallest integer for which the following statement holds: for every metric space M, M is L^1-embeddable in R^k iff every $c(k)$-sized subspace of M is L^1-embeddable in R^k. A special case of a theorem of Menger (see Blumenthal, *Theory and Applications of Distance Geometry*, Oxford Univ. Press, (1953), p.94) says that $c(1)$ exists and equals 4. We show that $c(2)$ exists and satisfies $6 \leq c(2) \leq 11$. The proof entails an $O(n^3)$-time algorithm for L^1-embedding an n-point metric space M into the plane when such an embedding exists. We conjecture that $c(k)$ exists for all $k \geq 1$. Finally, we discuss a lower bound for $c(k)$.

Running head: L^1-EMBEDDABILITY OF METRIC SPACES

Mail all correspondance to:

Seth M. Malitz
Dept. Computer [redacted] Science
Univ. of Mass.
Amherst, MA 01003

1 Introduction

1.1 On L^2-Embeddability of Metric Spaces

Let M be a metric space, and R^k denote Cartesian k-space. Say M is *L^q-embeddable in R^k* if there is a distance-preserving map of M into R^k with the L^q-metric.

In 1928, Menger [6] demonstrated the following bounded compactness theorem for L^2-embeddability in R^k.

Theorem 1.1 *Fix any positive integer k. An arbitrary metric space M is L^2-embeddable in R^k iff every $(k+3)$-point subspace of M is L^2-embeddable in R^k.*

A nice presentation of this result can be found in Blumenthal [5], along with bounded compactness theorems for other host metric spaces. A key observation in the proof of Menger's theorem is the following fact: given any set of $t+1$ points in R^t not L^2-embeddable in R^{t-1}, and any set of L^2-distances from these points, there is at most one point in R^t satisfying these distances.

1.2 On L^1-Embeddability of Metric Spaces

It is natural to ask if analogs of Menger's Theorem hold for other L^q-metrics on R^k besides L^2. We are convinced they do, but have not seen them discussed in the literature. For $1 < q < \infty$, we suspect a proof would follow rather closely the proof of Menger's Theorem since given any set of $t+1$ points in R^t not L^q-embeddable in R^{t-1}, and any set of L^q-distances from these points, there is at most one point in R^t satisfying these distances. However, the cases $q = 1$ and $q = \infty$ are different because for any finite set of points in R^t, there is a set of L^1 or L^∞-distances (respectively) from these points, such that infinitely many points in R^t satisfy these distances.

The initial goal of this paper was to prove a bounded compactness theorem for L^1-embeddability of metric spaces in R^k, for any fixed $k \geq 2$. (For $k = 1$, the bounded compactness theorem is just a special case of Menger's Theorem since the L^1- and L^2-metrics are identical on R^1.) Although we fell shy of this goal, we were successful for the case $k = 2$. In Section 2, we prove that an arbitrary metric space M is L^1-embeddable in R^2 iff every 11-point subspace of M is L^1-embeddable in R^2. Our suspicions are that the number 11 can be replaced by 6, but we have not been able to prove this.

Since the mapping $f : R^2 \rightarrow R^2$ which rotates the plane 45 degrees and then shrinks it by a factor $1/\sqrt{2}$ is an isometry from (R^2, L^1) to (R^2, L^∞), our bounded compactness theorem for L^1-embeddability in the plane is also a bounded compactness theorem for L^∞-embeddability in the plane.

Let $c(k)$ be the smallest integer for which the following statement holds: for every metric space M, M is L^1-embeddable in R^k iff every $c(k)$-sized subspace of M is L^1-embeddable in R^k. We conjecture that $c(k)$ exists for all $k \geq 1$, and that $c(k) = 2k + 2$. In an earlier version of this paper, we showed $c(2) \geq 6$, and $c(k) \geq 2k + 1$ for all $k \geq 1$. Recently, Jim Schmerle [7] gave a different and more sophisticated argument to show $c(k) \geq 2k + 2$ for all $k \geq 1$.

As stated above, we have not been able to establish the existence of $c(k)$ for any $k \geq 3$. Our attempts to extend the ideas we used in the case $k = 2$ seem to get horribly complicated in higher dimensions. In Section 3, we describe some even simpler-sounding problems that we also cannot solve. We hope that some reader will give them a try.

1.3 Algorithmic Aspect of L^1-Embeddability

The *real sequential random access machine* (RSRAM) is a device which can manipulate real numbers and perform simple arithmetic operations on them in unit time with infinite precision. It happens that our *proof* of the existence of $c(2)$ translates into an $O(n^3)$-time algorithm (for the RSRAM) to detect whether or not an arbitrary n-point metric space is L^1-embeddable in the plane, and construct such an embedding when one exists. (We have emphasized the word *proof* because as far as we can tell, the mere existence of $c(2)$ is only enough to guarantee an algorithm for *detection*, not construction, of an L^1-embedding.)

Let k be a fixed positive integer. Clearly, the existence of $c(k)$ would imply a polynomial time algorithm for detecting the L^1-embeddability of finite metric spaces in R^k. Such an algorithm would be of interest since the L^1-embeddability problem for unrestricted k is known to be NP-complete (see Avis and Deza [3]).

2 L^1-Embeddability in the Plane

In this section, we prove

Theorem 2.1 *An arbitrary metric space M is L^1-embeddable in the plane iff every 11-point subspace of M is L^1-embeddable in the plane.*

To start the proof, let $M = (W, d)$ be an arbitrary metric space. Let (*) denote the assumption that every 11-point subspace of M is L^1-embeddable in the plane. In what follows, we show that (*) implies that every *finite* subspace of M is L^1-embeddable in the plane. Hence by the Compactness Theorem of Logic, M itself is L^1-embeddable in the plane.

2.1 Preliminary Definitions

Let $u = (x_u, y_u)$ and $v = (x_v, y_v)$ be two points in the plane. The *relative location* of u w.r.t. v is described as follows. Write $u \nearrow v$ if v is up and to the right of u. Say $u \nearrow v$ is *strict* if $x_u < x_v$ and $y_u < y_v$. Similarly define $u \searrow v$ and *strict* $u \searrow v$. If A and B are sets of points in the plane, write $A \nearrow B$ to mean all points of B lie above and to the right of all points in A. Similarly define $A \searrow B$.

Given a finite set V, a function $g : V \to R^2$ is called an *exact location* of V. We often write $g(V)$ to indicate that g is an exact location acting on the set V. If X is any subset of V, we write $g(V)|X$ to denote the restriction of $g(V)$ to X. Two exact locations are considered to be the *same* if one is merely a translate of the other. If $A, B \subseteq V$, then we write $g[A \nearrow B](V)$ to mean that the exact location $g(V)$ satisfies the relative location $A \nearrow B$.

Three points u, v, w in the plane form a *3-chain* with v as *center* if $u \nearrow v \nearrow w$ or $u \searrow v \searrow w$. Clearly, the L^1-distances associated with a 3-chain satisfy the *triangle equality.* Three points u, v, w in an abstract metric space $M = (W, d)$ form a *3-chain* with v as *center* if $d(u, v) + d(v, w) = d(u, w)$.

An L^1-embedding of a metric space M in the plane is called a *planar L^1-embedding* of M. Given a planar L^1-embedding of M, there is a smallest rectangle (drawn parallel to the coordinate axes) that circumscribes it. This rectangle is called the *bounding rectangle* of the embedding.

2.2 Points on the Bounding Rectangle

We begin with a lemma.

Lemma 2.1 *Let $M = (W, d)$ be a metric space for which (*) holds. Let $N = (V, d)$ be any finite subspace of M. Then there is a subspace B of N such that: (1) each point $z \in N - B$ is the center of a 3-chain involving two points of B, and (2) B does not contain a 3-chain.*

Proof Say a subset Y of N is a *bounding set* if for every point $z \in N - Y$, z is the center of a 3-chain involving two points of Y. We show that N has a bounding set B that does not contain a 3-chain.

To do this, we construct a decreasing sequence of bounding sets $Y_0 \supset Y_1 \supset \cdots \supset Y_k$, where Y_k does not contain a 3-chain. Let $Y_0 = M$. It is vacuously true that Y_0 is a bounding set. Now suppose we have obtained the bounding set Y_i, and Y_i contains a 3-chain, say $\{u, v, w\}$ with center v. Let $Y_{i+1} = Y_i - \{v\}$.

We claim Y_{i+1} is a bounding set. For contradiction, suppose not. Then since Y_i is a bounding set, there is a point $x \in N - Y_i$ such that all 3-chains involving two points of Y_i and x as center must contain v. Let $y \in Y_i$ be such that $\{v, x, y\}$ is a 3-chain with x as center. (Notice y may equal u or w.) Let $S = \{u, v, w, x\} \cup \{y\} \subseteq Y_i$. We claim that (S, d) is not planar L^1-embeddable. For if there were such an embedding, the points v, x, y would form a 3-chain in the plane with x as center, the points u, v, w would form a 3-chain in the plane with v as center, and thus (a quick drawing helps here) x would also be the center of a 3-chain in S not involving v. This contradicts the third sentence of this paragraph. Thus (S, d) is not planar L^1-embeddable. But this violates (*). Thus Y_{i+1} is a bounding set.

To finish the proof, we simply let $B = Y_k$. (Notice $|B| \geq 2$.) ■

From now on, we shall always assume (*) holds for M, and thus for any finite subspace $N = (V, d)$ of M. As remarked above, $|B| \geq 2$. We now claim $|B| \leq 4$. For contradiction, suppose not. By (*), every 5-point subspace of B would planar L^1-embeddable. But any 5 points in the plane contain a 3-chain. Thus B would contain a 3-chain, contradicting the definition of B. Hence $|B| \leq 4$.

It is clear from the definition of B, that in any planar L^1-embedding of N (assuming one exists), all points of B must lie on the bounding rectangle, and every edge of the bounding rectangle is incident to a point of B.

Suppose $|B| = 4$. Let $g(B)$ be any planar L^1-embedding of B. If there exists a point r in the plane, $r \notin B$, such that a translated coordinate system centered at r puts each point of $g(B)$ in a different quadrant (see Figure 1), then we can extend $g(B)$ to a planar L^1-embedding of N. This follows from two observations that are not difficult to see: (1) g is the unique planar L^1-embedding of B up to translation, reflection, and relabeling of axes; (2) every point of R^2 inside or on the rectangle induced by $g(B)$ satisfies a unique set of L^1-distances from the points of B. By (*), $B \cup \{v, w\}$ is planar L^1-embeddable for all $v, w \in N - B$. By (1) and (2), the planar L^1-embeddability of $B \cup \{v, w\}$ for all $v, w \in N - B$ implies the planar L^1-embeddability of N.

2.3 A Small Extension of N

If $|B| \neq 4$ or B has a planar L^1-embedding not of the above form, then significantly more work is required to demonstrate the planar L^1-embeddability of $N = (V, d)$. To do this, we begin by extending N to a new metric space $N_p = (V \cup \{p\}, d)$. Later we show how to construct a planar L^1-embedding for this extended metric space using p as a convenient reference point in the plane. We define N_p in accordance with the following cases.

Case (i): $|B| = 2$. Arbitrarily select p to be either point of B. Thus $N_p = N$.

Case (ii): $|B| = 4$. To be distinct from the case already considered in Section 2.2, B must have a planar L^1-embedding $g(B)$ like that shown in Figure 2. Now there is a point r in the plane such that a translated coordinate system centered at r puts two points x,y of B in one quadrant and the other two points u,v in the diagonally opposite quadrant. It is readily seen that all planar L^1-embeddings of B have this form (x, y in one quadrant, u, v in the diagonally opposite quadrant.) By (*), $B \cup \{z\}$ is planar L^1-embeddable for all $z \in N - B$.

Consider any particular $z \in N - B$. It is easy to see that the L^1-distance between z and the corner point p (see Figure 2) is constant (and say equals δ_{pz}) over all L^1-embeddings of $B \cup \{z\}$. To obtain N_p, extend the domain of d from pairs in V to pairs in $V \cup \{p\}$ by setting $d(p, z) = \delta_{pz}$ for each $z \in N - B$.

Case (iii): $|B| = 3$. In this case, there are two distinct relative locations Π_1 and Π_2 of B (where Π_1 cannot be obtained from Π_2 by reflection or relabeling of axes) such that B has a planar L^1-embedding satisfying Π_1 and another satisfying Π_2. These relative locations are depicted in Figures 3(a) and 3(b). In either case, there is exactly one point of B at the corner of the induced rectangle.

By (*), every 11-point subspace of N containing B is planar L^1-embeddable. Each of these planar L^1-embeddings satisfies either Π_1 or Π_2 on B. Thus every 7-point subspace of N containing B has a planar L^1-embedding that satisfies Π_1 on B, or every 7-point subspace of N containing B has a planar L^1-embedding that satisfies Π_2 on B. For if this were not the case, then there would be 7-point subspace $S_1 = B \cup \{s, t, u, v\} \subseteq N$ for which every planar L^1-embedding satisfies Π_1 on B, and another 7-point subspace $S_2 = B \cup \{w, x, y, z\} \subseteq N$ for which every planar L^1-embedding of S_1 satisfies Π_2 on B, and thus $S_1 \cup S_2$ would be an 11-point subspace of N with no planar L^1-embedding at all, thus violating (*).

Without loss of generality, suppose every 7-point subspace of N containing B satisfies Π_1 on B. To construct the extension N_p, we let $p = a$ (see Figure 3(a)). Thus $N_p = N$.

2.4 A Graph Construction

Let (†) denote the assumption that every 5-point subspace $(\{p, u, v, w, x\}, d)$ of N_p is planar L^1-embeddable with $p \nearrow \{u, v, w, x\}$. Notice that from the way N_p is defined, (*) implies (†). In the subsections that follow, we will show how to construct a planar L^1-embedding of N_p with $p \nearrow N_p - \{p\}$. The construction is based on a certain graph whose vertices are the points of $N_p - \{p\}$.

First, define an equivalence relation $\overset{*}{\sim}$ on the points of $N_p - \{p\}$ as follows. Say $u \sim v$ iff there is a planar L^1-embedding of $(\{p, u, v\}, d)$ such that $p \nearrow \{u, v\}$ and strict $u \searrow v$. (Notice that for all such embeddings, the exact location of u w.r.t. v is is always the same.)

Let $\stackrel{*}{\sim}$ be the reflexive, transitive closure of $\sim$.

Let $V_1, \ldots, V_k$ be the equivalence classes of $\stackrel{*}{\sim}$. For each i, let E_i be those vertex pairs in V_i that are directly related by $\sim$, and let G_i be the graph (V_i, E_i).

2.5 Creating a Planar L^1-Embedding for N_p

Let T_i be any breadth-first search tree in G_i rooted at some node r_i. Suppose u,v,w is any path of length 2 in T_i. By (†) and the definition of $\sim$, there exists a planar L^1-embedding of $(\{p, u, v, w\}, d)$ satisfying $p \nearrow \{u, v, w\}$ and strict $u \searrow v$. Notice that in all such embeddings, the exact location of w w.r.t. v is always the same.

Let (r_i, s_i) be any edge of T_i. Let $f(\{p, r_i, s_i\})$ be any planar L^1-embedding of $(\{p, r_i, s_i\}, d)$ satisfying $p \nearrow \{r_i, s_i\}$ and strict $r_i \searrow s_i$. Let $f(\{r_i, s_i\})$ be the restriction of $f(\{p, r_i, s_i\})$ to $\{r_i, s_i\}$. We extend $f(\{r_i, s_i\})$ to an exact location of T_i inductively as follows. Let T_i' be any subtree of T_i that contains the edge (r_i, s_i). Suppose the exact location $f(T_i')$ has already been constructed. Let w_i be any vertex of $T_i - T_i'$ such that (w_i, v_i) is an edge of T_i for some vertex v_i in T_i'. Let (u_i, v_i) be an edge of T_i'. By (†), $(\{p, u_i, v_i, w_i\}, d)$ is planar L^1-embeddable with $p \nearrow \{u_i, v_i, w_i\}$ and strict $u_i \searrow v_i$. By the preceding paragraph, all such embeddings place w_i in the same exact location w.r.t. v_i. Extend $f(T_i')$ in accordance with this unique placement of w_i w.r.t. v_i. Continue in this manner, expanding the current tree one vertex at a time until an exact location for all of T_i is obtained. Call the resulting exact location $f(T_i)$. Notice that $f(T_i)$ is independent of the edge we started with—any starting edge for T_i will yield the same exact location $f(T_i)$.

Clearly, the above construction can be applied to any connected acyclic subgraph H of G_i.

Definition 2.1 *Let (u, v) be an edge in in the connected acyclic subgraph H of G_i. Then $f[u \searrow v](H)$ shall denote the exact location for H (satisfying strict $u \searrow v$) obtained by the method described above.*

A cycle C in G_i will be called *minimal* if the subgraph of G_i induced by the vertices of C equals C, i.e., C has no chords. Similarly, a path P in G_i will be called *minimal* if the subgraph of G_i induced by the vertices of P equals P.

Lemma 2.2 *Let $P = u_1, u_2, \ldots, u_t$ be a minimal path in G_i. Without loss of generality, $f[u_1 \searrow u_2](P)$ satisfies the relative locations depicted in Figure 4.*

Proof. Let P_1 be the length-3 subpath u_1, u_2, u_3, u_4 of P. Since (†) and $u_1 \sim u_2$ hold, $(P_1 \cup \{p\}, d)$ has a planar L^1-embedding with $p \nearrow P_1$ and strict $u_1 \searrow u_2$. By definition, $f[u_1 \searrow u_2](P_1)$ is the restriction of this embedding to P_1. Since P_1 is a minimal path, $f[u_1 \searrow u_2](P_1)$ satisfies $u_1 \nearrow u_3$ or $u_1 \swarrow u_3$. Without loss of generality, assume $u_1 \nearrow u_3$. Thus $f[u_1 \searrow u_2](P_1)$ satisfies the relative locations depicted in Figure 4.

Similarly, let P_2 be the length-3 subpath u_2, u_3, u_4, u_5. Since (†) and $u_2 \sim u_3$ hold, $(P_2 \cup \{p\}, d)$ is planar L^1-embeddable with $p \nearrow P_2$ and $u_3 \searrow u_2$. Hence by the minimality of P_2, $f[u_3 \searrow u_2](P_2)$ satisfies the relative locations depicted in Figure 4.

By sliding this "window" of length 3 over the entire length of the minimal path P, we see that $f[u_1 \searrow u_2](P)$ satisfies the relative locations of Figure 4. ■

Lemma 2.3 *Let $P = u_1, u_2, \ldots, u_t$ be a minimal path in G_i. Suppose $f[u_1 \searrow u_2](P)$ satisfies the relative location shown in Figure 4. Let $f[u_1 \searrow u_2](P \cup \{p\})$ be an exact location of $P \cup \{p\}$ whose restriction to P is $f[u_1 \searrow u_2](P)$, and whose restriction to $\{p, u_1, u_2\}$ is a planar L^1-embedding of $(\{p, u_1, u_2\}, d)$ with $p \nearrow \{u_1, u_2\}$. Then $f[u_1 \searrow u_2](P \cup \{p\})$ is a planar L^1-embedding of $(P \cup \{p\}, d)$ satisfying $p \nearrow P$.*

Proof Given how $f[u_1 \searrow u_2](P)$ is constructed, clearly the L^1-distance between p and any u_i in $f[u_1 \searrow u_2](P \cup \{p\})$ is equal to $d(p, u_i)$. Also, given any two vertices u_i, u_k in P of distance at most 3 from each other, (†) implies that the L^1-distance between u_i and u_k in $f[u_1 \searrow u_2](P \cup \{p\})$ equals $d(u_i, u_k)$.

As an induction hypothesis, suppose that for any two vertices u_i, u_k in P of distance at most s (where $s \geq 3$), the L^1-distance between them in $f[u_1 \searrow u_2](P \cup \{p\})$ equals $d(u_i, u_k)$.

We now show the same holds for distance $s + 1$. Suppose u_i, u_k are any two vertices in P of distance $s + 1 \geq 4$. Glancing at Figure 4, there is a vertex u_j in P such that

$u_i \nearrow u_j \nearrow u_k$ is satisfied by $f[u_1 \searrow u_2](P \cup \{p\})$. By the induction hypothesis, the L^1-distance between u_i and u_j in $f[u_1 \searrow u_2](P \cup \{p\})$ equals $d(u_i, u_j)$, and the L^1-distance between u_j and u_k in $f[u_1 \searrow u_2](P \cup \{p\})$ equals $d(u_j, u_k)$. Thus, any planar L^1-embedding of $(\{p, u_i, u_j, u_k\}, d)$ with $p \nearrow \{u_i, u_j, u_k\}$ (such an embedding is guaranteed to exist by (†)) satisfies $p \nearrow u_i \nearrow u_j \nearrow u_k$. Hence the L^1-distance between u_i and u_k in $f[u_1 \searrow u_2](P \cup \{p\})$ equals $d(u_i, u_k)$. ■

Lemma 2.4 *Any minimal cycle C in G_i has length at most 4.*

Proof Suppose, for contradiction, that $C = u_1, u_2, \ldots, u_k, u_1$ is a minimal cycle in G_i of length at least 5. Let P be the path $u_1, u_2, \ldots, u_k$. Since C is minimal, every length 3 subpath of P is minimal. By Lemma 2.2, $f[u_1 \searrow u_2](P)$ satisfies the relative locations of Figure 4. By Lemma 2.3, $f[u_1 \searrow u_2](P \cup \{p\})$ is a planar L^1-embedding of $(P \cup \{p\}, d)$ satisfying $p \nearrow P$. But $u_1 \nearrow u_k$ in this embedding. Thus (u_1, u_k) is not an edge of G_i. But this violates the fact that C is a cycle in G_i. ■

Lemma 2.5 *$f(T_i)$ is a planar L^1-embedding of (T_i, d).*

Proof First, we show that $f(T_i)$ observes the correct L^1-distances between all vertex pairs x, y in T_i where y is a descendant of x. Since T_i is a breadth-first search tree in G_i, the path P in T_i connecting x and y is a minimal path in G_i. Clearly, $f(T_i)|P = f(P)$. By Lemma 2.3, $f(P)$ is a planar L^1-embedding of (P, d). Hence the L^1-distance between x and y in $f(T_i)$ is correct.

It remains to show that $f(T_i)$ observes the correct L^1-distances between all vertex pairs x, y in T_i where neither x nor y is a descendant of the other. Let P be the path in T_i from x to y. Let Q be some *minimum* path (i.e., a path utilizing the smallest number of edges) in G_i from x to y. Since T_i is a breadth-first search tree, we may asume without loss of generality that the vertices and edges shared by Q and T_i form an initial subpath of Q and a final subpath of Q. See Figure 5. Let C be the indicated cycle.

By Lemma 2.3, $f(Q)$ is a planar L^1-embedding of (Q, d). Hence the L^1-distance between x and y in $f(Q)$ is correct. What we want to show is that the L^1-distance between x and y

in $f(T_i)|P$ (which equals $f(P)$) is correct. To do this, we prove that $f(P)$ and $f(Q)$ place y in the same exact location w.r.t. x. By (†), it suffices to show that there is an exact location $g(C)$ for C that is *2-consistent*, i.e., for every length-2 subpath a, b, c of C, there holds $g(C)|(a,b,c) = f((a,b,c))$.

Claim : *For any cycle Γ in G_i, Γ has an exact location $g(\Gamma)$ that is 2-consistent.*

The proof is by induction on the length of Γ.

Suppose the length of Γ is at most 4. Then by (†), we know that $(\Gamma \cup \{p\}, d)$ has a planar L^1-embedding, call it $g(\Gamma \cup \{p\})$, with $p \nearrow \Gamma$. Let $g(\Gamma) = g(\Gamma \cup \{p\})|\Gamma$. Clearly, $f((a,b,c)) = g(\Gamma)|(a,b,c)$ for any length-2 subpath a, b, c in Γ.

Now suppose the claim is true for all cycles Γ in G_i of length $\leq k$, for some fixed $k \geq 4$.

Consider a cycle Γ in G_i of length $k+1$. We show that Γ has an exact location $g(\Gamma)$ that is 2-consistent. Since the length of Γ is greater than 4, it follows from Lemma 2.4, that Γ has a chord (s,t) in G_i. Let Γ_1 and Γ_2 be the two induced cycles (see Figure 6). By the induction hypothesis, Γ_1 has a 2-consistent exact location $g(\Gamma_1)$ with strict $s \searrow t$, and Γ_2 has a 2-consistent exact location $g(\Gamma_2)$ with strict $s \searrow t$. Let $g(\Gamma)$ be the exact location of Γ obtained by glueing together $g(\Gamma_1)$ and $g(\Gamma_2)$ at the vertices s and t. Let α be the vertex other than t that is adjacent to s on Γ_1, and let β be the vertex other than t that is adjacent to s on Γ_2. Specializing (†) to the 5-space $(\{p, s, t, \alpha, \beta\}, d)$, it must be that $g(\Gamma)|(\alpha, s, \beta) = f((\alpha, s, \beta))$. ■

Lemma 2.6 *Let T_i be a breadth-first tree in G_i with root vertex r_i, where r_i is chosen so that $d(p, r_i) = \min_{x \in G_i} d(p, x)$. Then the planar L^1-embedding $f(T_i)$ can be extended to a planar L^1-embedding $f(T_i \cup \{p\})$ satisfying $p \nearrow T_i$.*

Proof Consider the set $S = \{s_1, \ldots, s_l\}$ of vertices such that (r_i, s_j) is an edge of G_i for $j = 1, \ldots, l$. By (†), there is a planar L^1-embedding of $(\{p, r_i, s_{j_1}, s_{j_2}, s_{j_3}\}, d)$ with $p \nearrow \{r_i, s_{j_1}, s_{j_2}, s_{j_3}\}$ for every 3-element subset $\{s_{j_1}, s_{j_2}, s_{j_3}\}$ of S. It follows that $(S \cup \{p\}, d)$ has a planar L^1-embedding, call it $f(S \cup \{p\})$ with $p \nearrow S$. Let $f(T_i \cup \{p\})$ be an exact placement of $T_i \cup \{p\}$ whose restriction to T_i is $f(T_i)$, and whose restriction to $S \cup \{p\}$ is $f(S \cup \{p\})$. Observe that $f(T_i \cup \{p\})$ satisfies $p \nearrow T_i$.

Now let v be any vertex in G_i. Let P be the path in T_i from r_i to v. Since P is a minimal path in G_i, Lemma 2.3 tells us that the L^1-distance between p and v in $f(T_i \cup \{p\})$ is equal to $d(p, v)$. ■

Consider the planar L^1-embedding $f(T_i)$. Of those vertices possessing least y-value, let b_i be the vertex with least x-value. Of those vertices in $f(T_i)$ possessing least x-value, let l_i be the vertex with least y-value. Of those vertices possessing largest y-value, let t_i be the vertex with largest x-value. Of those vertices possessing largest x-value, let r_i be the point with largest y-value.

We now show that N_p is planar L^1-embeddable with $p \nearrow N_p - \{p\}$. Let $b_1, \ldots, b_k$ be a listing of the b_i in order of increasing distance from p. By (†), the 5-point space $(\{p, t_i, r_i, l_{i+1}, b_{i+1}\}, d)$ is planar L^1-embeddable with $p \nearrow \{t_i, r_i, l_{i+1}, b_{i+1}\}$ for each i. Since t_i and r_i are in a different equivalence class from l_{i+1} and b_{i+1}, it follows that any planar L^1-embedding of $(\{p, t_i, r_i, l_{i+1}, b_{i+1}\}, d)$ satisfying $p \nearrow \{t_i, r_i, l_{i+1}, b_{i+1}\}$ also satisfies $p \nearrow \{t_i, r_i\} \nearrow \{l_{i+1}, b_{i+1}\}$. Furthermore, any such planar L^1-embedding coincides with $f(T_i)|\{t_i, r_i\}$ and $f(T_{i+1})|\{l_{i+1}, b_{i+1}\}$.

Let $f(N_p)$ be any exact location of N_p such that:

- $p \nearrow N_p - \{p\}$;
- $f(N_p)|T_i = f(T_i)$;
- $f(N_p)|\{p, t_i, r_i, l_{i+1}, b_{i+1}\}$ is a planar L^1-embedding of $(\{p, t_i, r_i, l_{i+1}, b_{i+1}\}, d)$ for each i.

The claim is that $f(N_p)$ is a planar L^1-embedding of N_p. We already know from Lemma 2.6, that $f(N_p)|(T_i \cup \{p\})$ is a planar L^1-embedding for each i. So what has to be shown now is that for any two points $u \in T_i$, $v \in T_j$, $i < j$, the L^1-distance between u and v in $f(N_p)$ is $d(u, v)$. By (†) and the assumption $i < j$, $(\{p, u, v\}, d)$ is planar L^1-embeddable with $p \nearrow u \nearrow v$. Thus $d(u, v) = d(p, v) - d(p, u)$. We know that in $f(N_p)$, there holds $p \nearrow u \nearrow v$, the L^1-distance between p and u is $d(p, u)$, and the L^1-distance between p and v is $d(p, v)$. Therefore in $f(N_p)$, the L^1-distance between u and v is $d(u, v)$.

This completes the proof of Theorem 2.1. We have shown that assuming (*), any finite metric space $N = (V, d)$ is planar L^1-embeddable. Hence, any infinite metric space satisfying (*) is also L^1-embeddable. ■

3 Higher Dimensions—Two Conjectures

Surprisingly, we have not been able to show that $c(k)$ exists for any $k \geq 3$. Our attempts to extend the methods of the last section seem to get wildly complicated in higher dimensions. The purpose of this section is to describe two related problems that appear to be simpler, but which we still cannot solve. A solution to either of these two problems might provide useful insights for demonstrating the existence of $c(k)$ for all k.

First we need a few definitions. Let $M = (W, d)$ be an arbitrary finite metric space. Let W consist of elements $p_1, p_2 \ldots, p_n$. Let G be a complete digraph with vertex set W each of whose arcs (p_i, p_j) is labeled with a k-vector σ_{ij} consisting only of +1's and −1's, and such that $\sigma_{ij} = -\sigma_{ji}$. Furthermore, assume the labeling is *transitive* in the following sense: if the lth coordinate of both σ_{ij} and σ_{ju} is +1, say, then the lth coordinate of σ_{iu} is +1. Similarly for −1. Call such an object G a *k-dimensional transitive plane assignment* (k-TPA).

Let f be an L^1-embedding of M into R^k. Let $f(p_i)[l]$ denote the lth coordinate of $f(p_i)$. Let $\sigma_{ij}[l]$ denote the lth coordinate of σ_{ij}. Say f *satisfies* G if for every $i, j \in \{1, \ldots, n\}$ with $i \neq j$, and for every $l \in \{1, \ldots k\}$, there holds $\sigma_{ij}[l] = +1$ implies $f(p_i)[l] \leq f(p_j)[l]$, and $\sigma_{ij}[l] = -1$ implies $f(p_i)[l] \geq f(p_j)[l]$.

Conjecture 3.1 *Fix any positive integer k. There is a constant $\epsilon(k)$ such that for any metric space M and any k-TPA G for M, M has an L^1-embedding in R^k that satisfies G iff every $\epsilon(k)$-sized subspace N of M has an L^1-embedding in R^k that satisfies G (restricting G to the points of N.)*

Given any mapping $g : W \rightarrow R^k$, and k-TPA G, let the *G-distance* between $g(p_i)$ and $g(p_j)$ be the quantity $[g(p_i) - g(p_j)] \cdot \sigma_{ij}$. In words, the G-distance between $g(p_i)$ and $g(p_j)$ is the least L^1-distance between $g(p_i)$ and a point on the hyperplane $(x_1, \ldots, x_k) \cdot \sigma_{ij} + g(p_j)$ (or equivalently, the least L^1-distance between $g(p_j)$) and a point on the hyperplane

$(x_1, \ldots, x_k) \cdot \sigma_{ji} + g(p_i)$. A *G-embedding* for M is a mapping $g : W \rightarrow R^k$ such that $[g(p_i) - g(p_j)] \cdot \sigma_{ij} = d(p_i, p_j)$ for all p_i and p_j.

Conjecture 3.2 *There is a constant $\gamma(k)$ depending only on k such that for any finite metric space $M = (W, d)$ and any k-TPA G for M, M has an G-embedding iff every $\gamma(k)$-sized subspace N of M has a G-embedding (restricting G to the points of N.)*

References

[1] P. Assouad, M.Deza, *Espaces metriques plongeables dans un hypercube: aspects combinatoires*, Annals of Discrete Math. 8 (1980) 197–210.

[2] P. Assouad, M. Deza, *Metric Subspaces of L^1*, Publications Mathematiques D'Orsay, Universite de Paris-Sud (1980).

[3] D. Avis, M. Deza, *The cut cone, L^1-embeddability, complexity and multicommodity flows*, to appear in Networks.

[4] L. M. Blumenthal, *Distance Geometries*, University of Missouri Studies, 13, no. 2 (1938).

[5] L. M. Blumenthal, *Theory and Applications of Distance Geometry*, Oxford Univ. Press, Clarendon, Oxford (1953).

[6] K. Menger, Untersuchungen uber allgemeine Metrik, *Mathematische Annalen*, 100 (1928), 75–163.

[7] James Schmerle, private communication. (1990)

Figure 1. A planar L^1-embedding of $B = (\{w, x, y, z\}, d)$ where each point of B lies in a different quadrant.

Figure 2. A planar L^1-embedding of $B = (\{u, v, x, y\}, d)$. Pairs u, v and x, y lie in diagonally opposing quadrants. The point p represents a corner of the bounding rectangle for B. The point z represents a member of $N - B$.

Figure 3. Two planar L^1-embeddings of $B = (\{a, b, c\}, d)$ satisfying different relative locations Π_1 (Figure 3(a)) and Π_2 (Figure 3(b)).

Figure 4. The relative locations satisfied by $f[u_1 \searrow u_2](P)$, where P is a minimal path in G_i.

Figure 5. The paths P and Q, and the induced cycle C.

Figure 6. The cycle Γ with chord (s, t), and the two induced cycles Γ_1 and Γ_2.

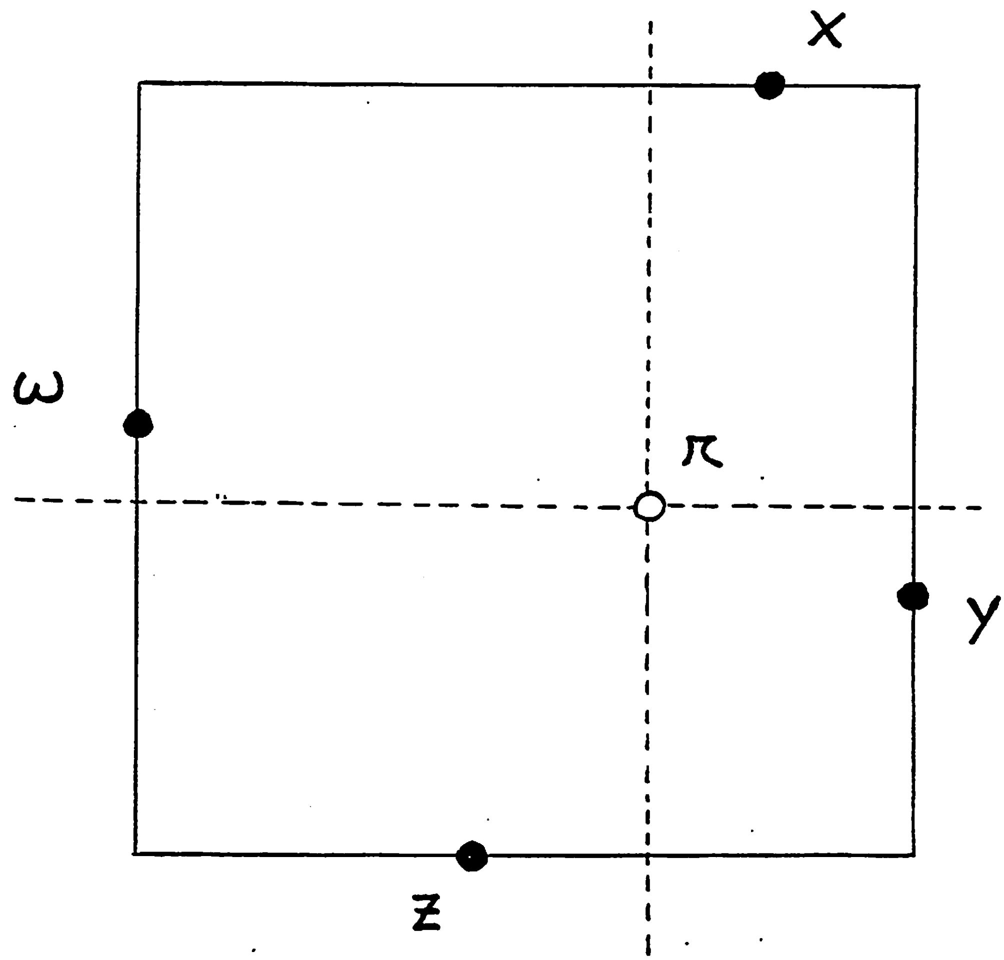
x
ω
π
y
z

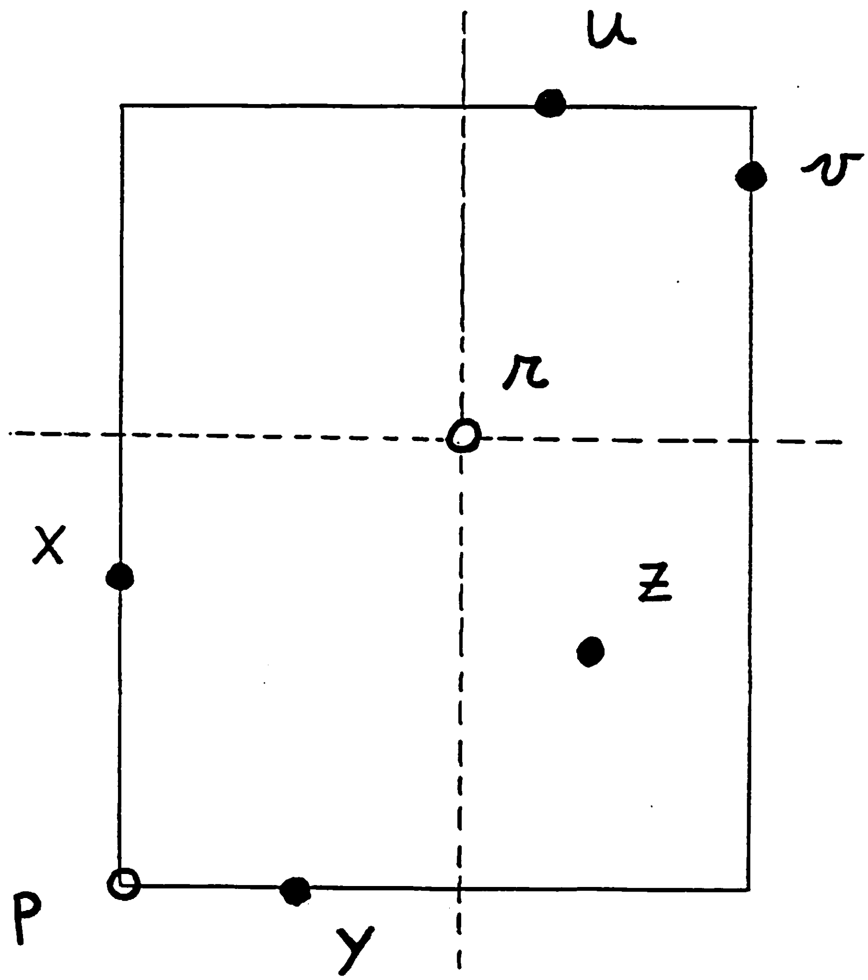

u
v
r
x
z
p
y

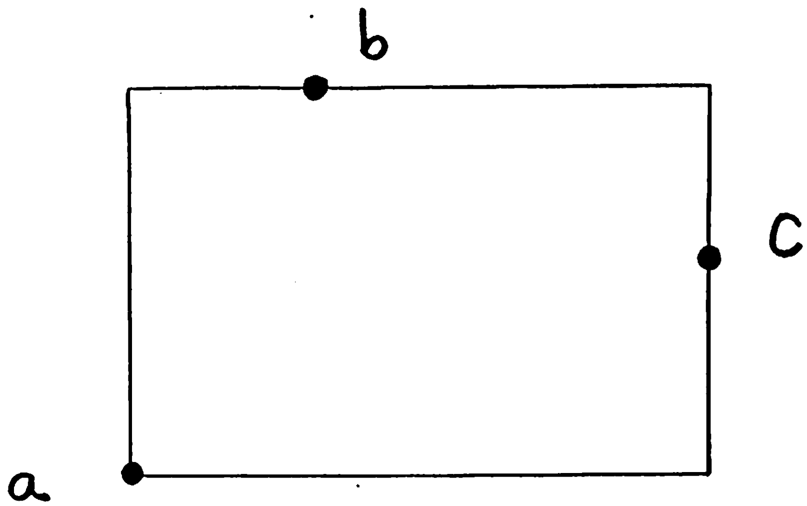

(a)

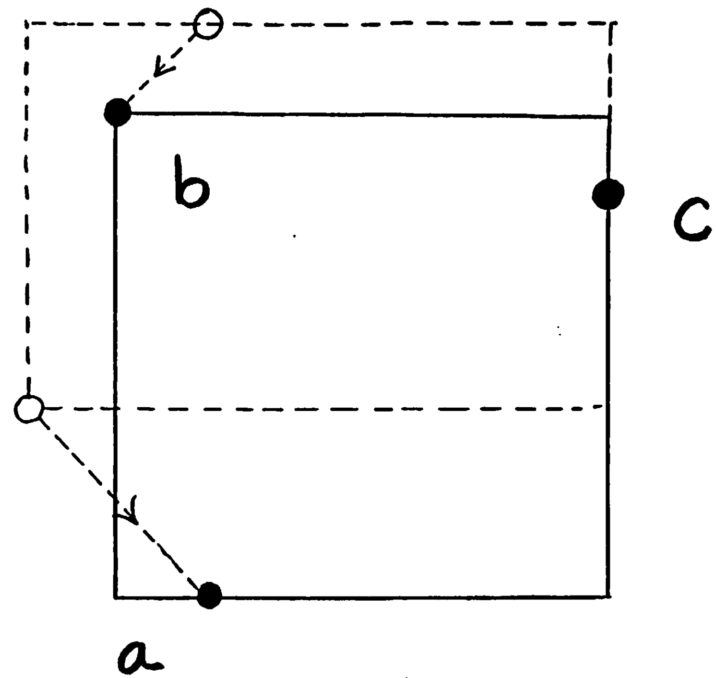

(b)

3

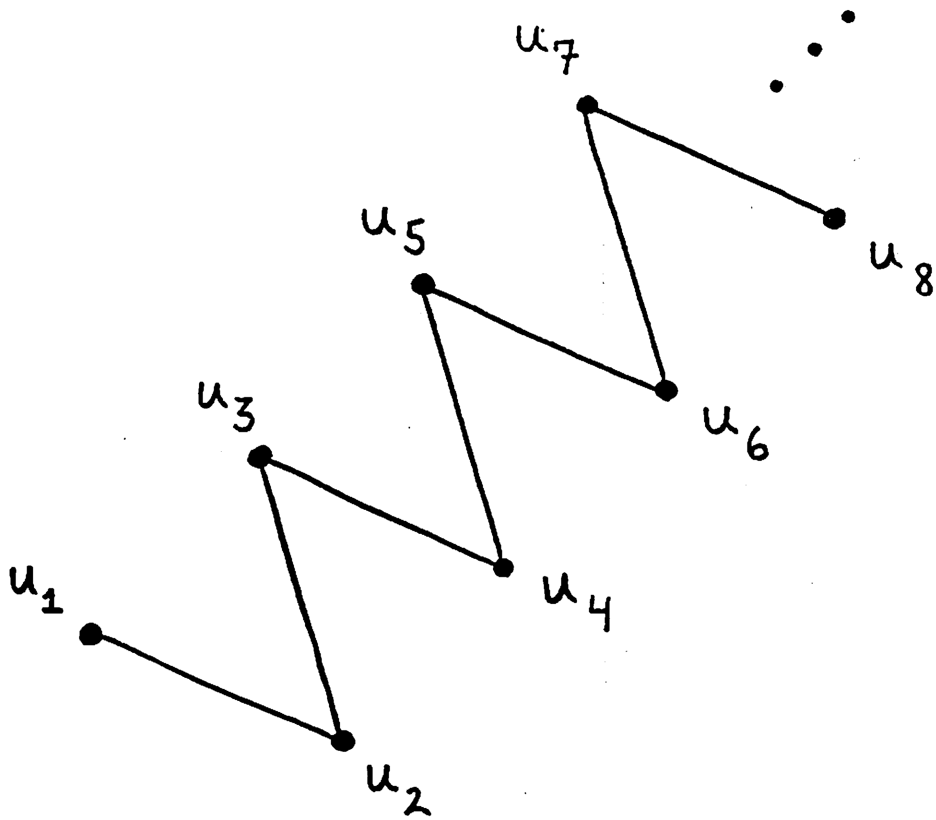
u_1
u_2
u_3
u_4
u_5
u_6
u_7
u_8

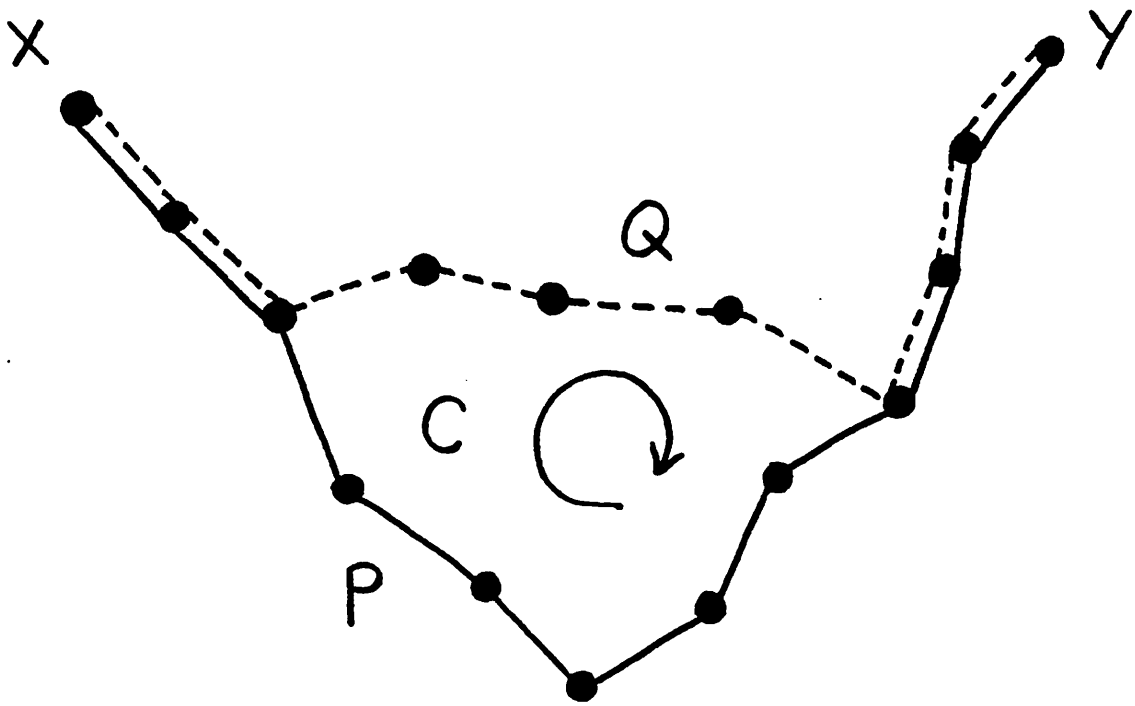
X
Y
Q
C
P

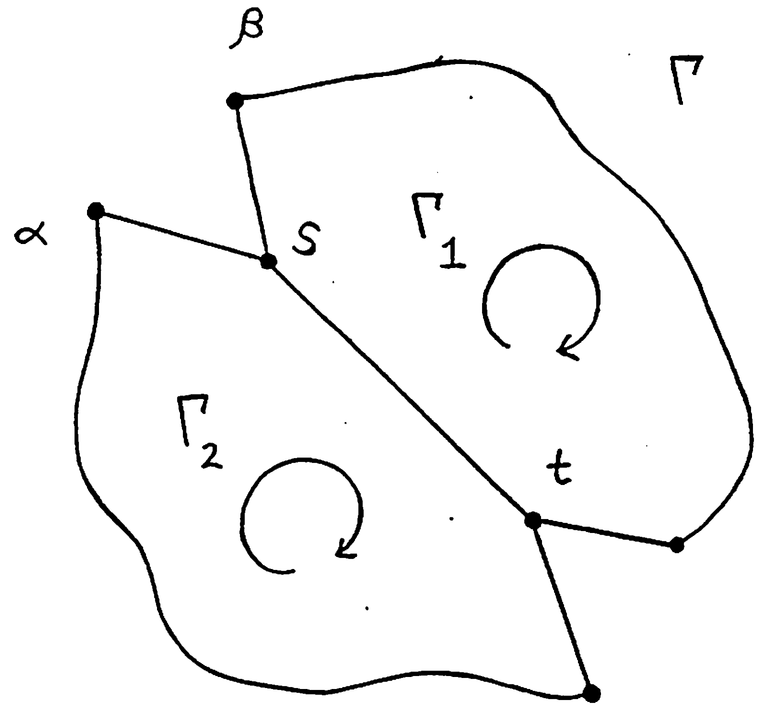
β
Γ
α
S
Γ_1
Γ_2
t