SIAM J. DISCRETE MATH.
Vol. 13, No. 2, pp. 268–279

DIRECTIONAL ROUTING VIA GENERALIZED *st*-NUMBERINGS*

FRED S. ANNEXSTEIN† AND KENNETH A. BERMAN†

Abstract. We present a mathematical model for network routing based on generating paths in a consistent direction. Our development is based on an algebraic and geometric framework for defining a directional coordinate system for real vector spaces. Our model, which generalizes graph *st*-numberings, is based on mapping the nodes of a network to points in multidimensional space and ensures that the paths generated in different directions from the same source are node-disjoint. Such directional embeddings encode the global disjoint path structure with very simple local information. We prove that all 3-connected graphs have 3-directional embeddings in the plane so that each node outside a set of extreme nodes has a neighbor in each of the three directional regions defined in the plane. We conjecture that the result generalizes to k-connected graphs. We also show that a directed acyclic graph (dag) that is k-connected to a set of sinks has a k-directional embedding in $(k-1)$-space with the sink set as the extreme nodes.

Key words. graph connectivity, network routing, *st*-numbering, matchings

AMS subject classifications. 68R10, 05C40, 68R10

PII. S0895480198333290

1. Introduction.

A fundamental problem in network routing is the generation of communication paths from a set of source nodes Y to a set of sink nodes X. Routing schemes for communication networks are often implemented using local tables that logically associate each destination address x with a parent link on which to forward messages to x (see [3, 6]). The set of parent links forms a directed *sink tree* to the sink node x. More generally, given a set of sink nodes X, we define a *sink tree to* X as a set of parent links, with one parent for each $u \notin X$, forming a union of directed in-trees, where each in-tree has a single sink node from X. When used in routing applications, the tables associated with such sink trees suffer from the fact that they do not provide a full representation of the entire network and therefore are vulnerable to faults and other dynamic network changes. One approach to the problem of fault-tolerant routing is the study of collections of *independent* sink trees and closely related graph *st*-numberings [10, 2]. A pair of sink trees to X is *independent* if it has the property that for each node $u \notin X$, the pair of tree-paths from u to X is internally node-disjoint. A standard way to obtain a pair of independent sink trees is via an *st*-numbering of a biconnected graph. Let $G = (V, E)$ be a graph with distinguished vertices s, t. An *st-numbering* of G is a mapping f of the vertices into the real line, $f : V \to R$, such that for all $v \in V - \{s, t\}$, there are neighbors $u, w \in N(v)$ such that $f(u) < f(v) < f(w)$. A 2-connected graph has an *st*-numbering for any vertex pair s, t [11, 8]. In general, collections of $k > 2$ sink trees are independent if they are pairwise independent. There is a long-standing open conjecture of Frank (see [13]) that states that all k-connected graphs have, for each vertex v, k-independent sink trees to v. One goal of this work is to provide a generalized graph *st*-numbering in an attempt to characterize graphs having k-independent sink trees.

*Received by the editors January 29, 1998; accepted for publication (in revised form) September 22, 1999; published electronically April 6, 2000.
http://www.siam.org/journals/sidma/13-2/33329.html

†Department of ECE & Computer Science, University of Cincinnati, Cincinnati, OH 45221 (fred.annexstein@uc.edu, ken.berman@uc.edu). This research was supported in part by NSF grant CCR-9877139 and by an OBR investment grant.

An st-numbering provides a stronger model for routing applications than that which is provided by the (induced) independent sink trees in the following sense: For each pair of vertices $u, v \notin X = \{s, t\}$, there exists a pair of tree-paths from distinct sink trees, one from u to X and the other from v to X, that is internally node-disjoint. We say that a pair of sink trees with this property is *strongly independent*. A collection of $k > 2$ sink trees is *strongly independent* if for each set Y of k vertices $\{y_1, \ldots, y_k : y_i \notin X\}$ there exists a matching M of Y to the k trees such that the collection of induced tree-paths (each routing a distinct y_i to X) is internally node-disjoint. We strengthen Frank's conjecture by conjecturing that every k-connected graph has k strongly independent sink trees, and we prove our conjecture for $k = 3$ using a generalized st-numbering that yields the matching M of Y to the three sink trees.

Our generalization of graph st-numbering is based on an algebraic and geometric framework for defining directions in real vector spaces. Using this framework we propose a mathematical model for network routing based on *directional graph embeddings with compass*. We show that such embeddings naturally induce collections of strongly independent sink trees and thus encode the global disjoint path structure with very simple local information. A k-directional embedding of a graph G involves mapping the vertices $V(G)$ to $(k-1)$-space so that each vertex, outside a set of extreme vertices, has a neighbor in each of k distinct directions. Routing in the directional embedding model is based on generating routing paths that are *monodirectional*, i.e., the paths follow a consistent direction. The mechanisms of the model ensure that the monodirectional paths generated in different directions from the same source are node-disjoint.

The compass associated with our embeddings yields directions (or equivalently, directional regions relative to any fixed location) that possess a fundamental matching property of combinatorial interest. Namely, any set of k distinct points in $(k-1)$-dimensional space can be matched to the set of k directions defined by the compass. The k points are matched in the sense that there exists a fixed location point s such that each of the original k points lies in a distinct directional region with respect to s. Furthermore, we characterize the set of all such points that satisfy this property, and we provide the conditions under which the associated matching is unique.

In this paper we apply these developments to prove that all 3-connected graphs have 3-directional embeddings in the plane. We conjecture that the result generalizes to the case of k-connected graphs, for each $k > 3$. Supporting this conjecture, we prove that a directed acyclic graph (dag) that is k-connected to a set of sinks has a k-directional embedding in $(k-1)$-space. A proof of our conjecture would affirm the conjecture of Frank.

Comparisons with previous work. The concept of graph st-numberings has been generalized to geometric embedding of graphs in Euclidean space in Linial, Lovász, and Widgerson [12]. Cheriyan and Reif [5] later extended the results to digraphs. This generalization of st-numberings originated by appealing to potential functions in physics (e.g., rubber bands or electricity), leading to the notion of a *convex embedding in general position*. A graph has a convex X-embedding f if f is a mapping of the vertices to points in real space $f : V \to R^{|X|-1}$ such that, for each $v \in V - X, f(v)$ is in the convex hull of $f(N(v))$. A convex embedding is in general position if $f(V)$ is in general position. Convex embeddings in general position characterize k-connected graphs, analogous to the st-numbering characterization of 2-connected graphs. Unfortunately, from the point of view of path generation, convex

X-embeddings for $|X| \geq 3$ neither seem to provide a model from which to extract structural information concerning disjoint paths to X nor seem to provide a means to generate independent sink trees.

Our conjecture that k-connected graphs have k-directional embeddings in R^{k-1} is equivalent to showing that convex embeddings in general position can be transformed to nondegenerate directional embeddings (see section 4). Our result that 3-connected graphs have 3-directional embeddings implies that 3-connected graphs have 3-independent spanning trees, which was previously shown by Zehavi and Itai [15] and Cheriyan and Maheshwari [4].

2. A directional compass. The basic framework for defining directions associated with a directional embedding is a natural one, obtained by defining a *k-directional compass* on the geometry of real vector spaces. A compass is given by a *direction function* $d : R^{k-1} \times R^{k-1} \to D$, where $D = \{1, 2, \ldots, k\}$ is a set of k directions, and has the following two properties: (1) A compass is *translatable* from the origin $\bar{\mathbf{0}}$, i.e., the direction of vector y relative to a reference vector x is $d(x, y) = d(\bar{\mathbf{0}}, y - x)$, for every pair of points x, y; and (2) a compass has *directional transitivity*, i.e., if $d(x, y) = d(y, z)$, then $d(x, y) = d(x, z)$, for every set of points x, y, z. From these two properties it follows immediately that for graphs embedded in k-space, monodirectional paths, i.e., paths that consistently follow the same direction, originating from the same vertex and following different directions, are internally node-disjoint. As a simple illustrative example, consider a graph embedded on the real line along with a 2-directional compass (on the real line) given by the positive and negative directions. It is clear that positive directed paths are disjoint from negative directed paths, when originating from the same node. It is this notion of directions in graphs that we wish to generalize to higher dimensions.

The generalized compass that we will apply enjoys a useful directional splitting property that is also enjoyed by the 2-directional compass on the real line. This splitting property of a compass is particularly useful for algorithmic constructions of directional embeddings.

DEFINITION. *A compass with k directions has the* k-directional splitting property *if, given any set of k distinct points $Y = \{y_1, y_2, \ldots, y_k\}$, there exists a reference point $s = s_Y$, called a* k-directional splitter point*, and an associated* directional permutation *π_Y, so that the point $y_{\pi_Y(i)}$ is in the ith direction relative to s, i.e., $d(s, y_{\pi_Y(i)}) = i$. A splitter point s is called a* strict-splitter *if each of the k points lies strictly in the interior of a directional region with reference to s, i.e., no point lies on a boundary shared by more than one region.*

The compass for the plane R^2 using the four natural directions N,S,E,W does not satisfy the 4-directional splitting property, as seen in Figure 1. We solve the problem of directional splitting by using a compass with three directions in R^2 and show that this extends to R^k using a compass with $k + 1$ directions.

We define a compass so that directional regions are associated with the space spanned by a given fixed set of vectors. To begin, let L be a fixed set of lines in R^k through the origin. Divide each line $l_i \in L$ at the origin into two half-lines and arbitrarily label one half-line as *ray-i* and the other half-line as *antiray-i*. We say that a particular labeling is a *convex bipartition labeling* of L if the set of chosen rays are *convex-spanning*, i.e., the entire space R^k is generated by positive (convex) linear combinations of vectors that lie on the rays. It follows that if the set of rays are convex-spanning, then the set of antirays are convex-spanning too. Clearly, in k-space, at least $k+1$ lines are required for there to exist a set of rays that is convex-spanning.

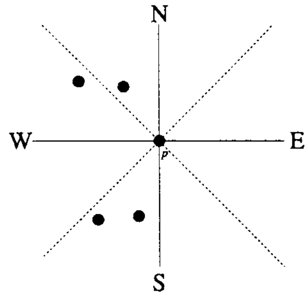

FIG. 1. *Four points that have no* 4*-directional splitter point. With respect to the reference point p shown, only three of four directional regions are covered. The dashed lines depict the boundaries of the directional regions. Further, no matter where p is positioned, the four fixed points cannot be located in four distinct directional regions relative to p.*

A set of points X in k-space is in *general position* if no d-element subset of X lies on a hyperplane of dimension $d-2$, for any $d \leq k+1$. We say that a set of lines L in k-space is in *general position* if the lines intersect at the origin and any set X of nonorigin points, chosen one per line, is in general position. The following shows that a set of $k+1$ lines in k-space that is in general position has a convex bipartition labeling, and this partition labeling is unique up to the naming of rays.

PROPOSITION 1. *Let L be a set of k lines in general position in $(k-1)$-space with one half-line from L marked as ray-*1*. Then there is a unique convex bipartition labeling of L.*

Proof. Arbitrarily order L, say, $\{l_1, l_2, \ldots, l_k\}$, where l_1 has the marked half-line. Label the marked half-line as ray-1, and label the other half-line as antiray-1. We now show that the remaining lines are uniquely labeled in a convex bipartitioning. Consider l_2 and the $(k-2)$-dimensional hyperplane H spanned by the remaining lines $\{l_3, l_4, \ldots, l_k\}$. Since H divides the original $(k-1)$-space into two half-spaces, it follows that a unique half-line of l_2 lies on the opposite side of H as the ray-1. Label this particular half-line as ray-2. Continuing in this way, each line $l_i \in L$ uniquely partitions into ray-i and antiray-i. The collection of rays must be convex-spanning since the lines are in general position. Any other labeling (in which ray-1 was fixed) could not be convex-spanning; thus the partition labeling is unique. ∎

DEFINITION. *Let L be a set of k lines in general position equipped with a convex bipartition labeling which marks half-lines as k rays and k antirays. The* compass $\mathcal{C}_k(L)$ *defined by L demarcates the* k-directional regions *relative to any fixed reference point. Each region is defined by the convex combinations of $k-1$ antirays chosen from the set of k antirays originating at the fixed reference point.*

The antirays thus mark the "boundary corners" of each region. Each of the k-directional regions is named by the index of the unique ray it contains (or equivalently, by the antiray it omits). The direction of a point that lies on the boundary (hyperplane) dividing more than one region is ambiguous, but by convention we consider it to have multiple directions, one for each region associated with the boundary.

There is, of course, a natural dual-compass demarcating *directional antiregions* defined by convex combinations of $k-1$ rays chosen from the set of k rays. Here the rays mark the boundary corners of directional antiregions, and each of the k-directional antiregions is named by the index of the unique antiray it contains. Antiregions are useful for identifying those points that have a fixed direction to a fixed reference point

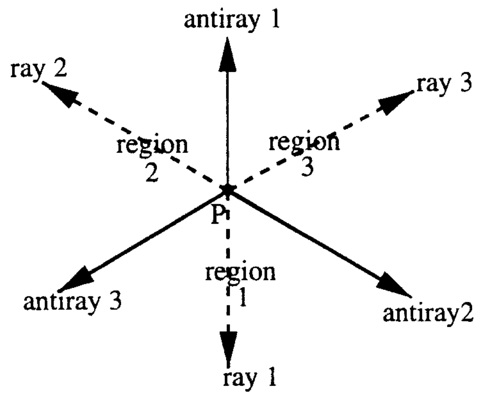

FIG. 2. *Compass with three regions in the plane from reference point p, obtained by a convex bipartitioning of three lines.*

and for helping to locate splitter points.

For example, in Figure 2 we have a compass $\mathcal{C}_3(L)$ in the plane, yielding directional regions that are symmetric. In this figure we see that region 1 is the set of convex combinations of points on antiray-2 and antiray-3. Dually, antiregion-1 (not labeled) is the set of convex combinations of points on ray-2 and ray-3 and defines precisely those points x for which point P lies in region 1, i.e., $d(x, P) = 1$. We can find a splitter point of a set of three points $\{A, B, C\}$, as shown in Figure 3, by considering the antiregions defined at each point. The splitter points are found in an intersection of three differently named antiregions. This intersection region associates each of the points $\{A, B, C\}$ to a different region. Note that in Figure 3 the set of all splitter points forms a convex region, and any interior splitter point divides $\{A, B, C\}$ into directional regions in the same way.

In the next section we show how to formalize this notion of matching points to directional regions, and we show how this extends to higher dimensions. We show that the set of splitter points always forms a convex region C and the set of strict-splitters is associated with the interior of C. Further, we show that there is a function that gives a weight to each point-direction pair such that maximum weight matchings correspond to the region C. The weighting is obtained by a special set of directional coordinates discussed in the next section.

3. A directional coordinate system and splitting lemma. For a given compass $C_k(L)$ in $(k-1)$-space, we now show that there exists a transformation from Euclidean coordinates to a *directional coordinate system.* This new system uses k-tuples in a manner where each fixed directional region (from the origin) is precisely those points whose ith component is maximum (dominant). The directional coordinatization is defined as follows. First, fix a collection $\{v_1, v_2, \ldots, v_k\}$ of k nontrivial vectors on the k rays defined by $C_k(L)$, precisely one per ray, such that $\sum v_i = \bar{0}$, the origin. Note that this can always be done since the rays are assumed to be convex-spanning.

To obtain directional coordinates, consider the linear transformation $\mathcal{T} : R^k \rightarrow R^{k-1}$ that maps each vector z in k-space with Euclidean coordinates $z = (z_1, z_2, \ldots, z_k)$ to the vector $-(z_1v_1 + z_2v_2 + \cdots + z_kv_k)$ in $(k-1)$-space. The nullspace of this linear mapping is the line through the "all-1s" vector $(1, 1, \ldots, 1)$. Thus for any two points

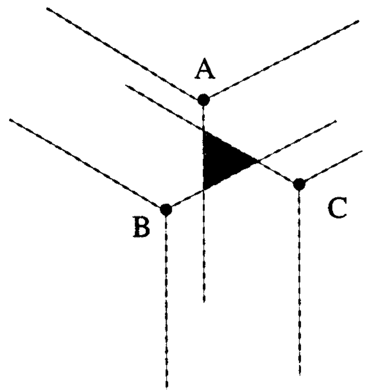

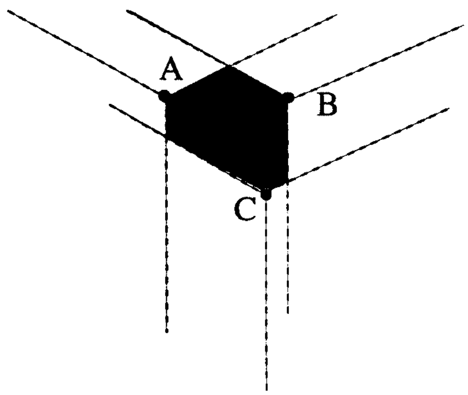

FIG. 3. *Two examples, each showing a set of three points $\{A, B, C\}$, the antiregions from each point, and the set of all splitter points shown as a shaded polytope. The directional (anti-)regions are defined by the compass in Figure 2. From each splitter point in the interior of the polytope, the points $\{A, B, C\}$ are matched to (lie in) distinct regions. In the top figure, each interior point matches A to region 2, B to region 3, and C to region 1. In the bottom figure, each interior point matches A to region 2, B to region 1, and C to region 3.*

$x_1, x_2 \in \mathcal{T}^{-1}(y)$ in the inverse transformation $\mathcal{T}^{-1}(y)$ of a point $y \in R^{k-1}$, we have that the differences between components are identical, and thus the minimum or maximum component of $\mathcal{T}^{-1}(y)$ is well defined. We will use angled brackets $\langle\rangle$ to represent the equivalence class and thereby denote the *directional coordinates*. So for example, we say that $\mathcal{T}^{-1}(y) = \langle y_1, y_2, \ldots, y_k\rangle$ is the set of directional coordinates associated with a point $y \in R^{k-1}$. As the following proposition shows, the maximum directional component of $\mathcal{T}^{-1}(y)$ identifies the directional region in $(k-1)$-space within which the point $y \in R^{k-1}$ lies, with respect to the origin.

First, we illustrate the concept of directional coordinates using the example of the compass from Figure 2. As noted, we first choose three vectors $\{v_1, v_2, v_3\}$ on these rays so that they sum to zero. We give the Euclidean coordinates of these vectors as follows: $v_1 = (0, -2)$ on ray-1, $v_2 = (-\sqrt{3}, 1)$ on ray-2, and $v_3 = (\sqrt{3}, 1)$ on ray-3. Now, say we are given a point q in the plane with Euclidean coordinates $q = (q_1, q_2)$. Under the inverse transformation $\mathcal{T}^{-1}(q)$ we obtain directional coordinates for q given by $q = \langle z_1, z_2, z_3\rangle$. By setting $q = -(z_1 v_1 + z_2 v_2 + z_3 v_3)$ and using the Euclidean coordinates for each v_i, we can solve for directional coordinates as follows: $z_1 = 0, z_2 = q_2/2 - q_1/2\sqrt{3}$, and $z_3 = z_2 + q_2$. So, for example, the point $q_0 = (0, 1)$

in Euclidean space has directional coordinates $q_0 = \langle 0, -\sqrt{3}/6, \sqrt{3}/6\rangle$. Since the maximum directional coordinate is the third, we have that q_0 lies in region 3.

PROPOSITION 2. *The linear transformation $\mathcal{T} : R^k \to R^{k-1}$ is an onto mapping with the property that the point $x \in R^{k-1}$ lies in the ith directional region ($1 \le i \le k$) relative to the origin iff the maximum directional-component of $\mathcal{T}^{-1}(x)$ is the ith component.*

Proof. The transformation $\mathcal{T}$ is an onto mapping since the collection of vectors $\{v_1, v_2, \dots, v_k\}$ spans $(k-1)$-space. By definition, a point $x \in R^{k-1}$ is in the ith directional region iff it is a convex combination of the antirays different from antiray-i, i.e., $x = \sum_{j \ne i} \lambda_j(-v_j)$, where each $\lambda_j \ge 0$ and each $-v_j$ is a vector on antiray-j. Hence, the directional coordinatization yields $\mathcal{T}^{-1}(x) = \langle -\lambda_1, -\lambda_2, \dots, -\lambda_{i-1}, 0, -\lambda_{i+1}, \dots, -\lambda_k\rangle$, and thus the ith component is maximum. $\square$

The following proposition is easily verified by applying Proposition 2.

PROPOSITION 3. *The directional compass is translatable and has directional transitivity, as defined in section* 2.

Within this directional coordinate system for $(k-1)$-space we can identify the problem of finding a directional permutation with the problem of finding a maximum perfect matching in an associated weighted bipartite graph or, equivalently, finding a maximum weight permutation. To wit, let $Y = \{y_1, y_2, \dots, y_k\}$ be a set of k points in Euclidean space R^{k-1}, and for each $1 \le i \le k$, let $\langle y_{i,1}, y_{i,2}, \dots, y_{i,k}\rangle$ denote (a representative of) $\mathcal{T}^{-1}(y_i)$. For π a permutation on A, we define $wt(\pi)$, the weight of π, as the sum $\sum y_{\pi(i),i}$.

We now prove the main theorem of this section.

THEOREM 1 (the k-direction splitting lemma). *Let A be a set of k points in R^{k-1}. The set of all k-directional splitter points of A forms a nonempty convex region. Further, a permutation π on A is a directional permutation associated with a splitter point iff π is a maximum weight permutation. Moreover, π is unique iff there exists a strict-splitter point of A.*

Proof. First, assume that π is a maximum weight permutation. We now show that we can find a splitter point x such that $y_{\pi(i)}$ is in the ith directional region relative to x. From Proposition 2 we know that such a point x must satisfy the following property: The vector $y_{\pi(i)} - x$ when coordinatized as $\mathcal{T}^{-1}(y_\pi(i) - x)$ must have its ith component as the maximum. Hence, choosing representatives $\langle x_1, x_2, \dots, x_k\rangle \in \mathcal{T}^{-1}(x)$ and $\langle y_{\pi(i),1}, y_{\pi(i),2}, \dots, y_{\pi(i),k}\rangle \in \mathcal{T}^{-1}(y_{\pi(i)})$, we have that the following inequalities must be satisfied (independent of the representative choices): for each $j \ne i$,

$$y_{\pi(i),i} - x_i \ge y_{\pi(i),j} - x_j.$$

These inequalities are simply a set of difference constraints of the form

$$x_i - x_j \le y_{\pi(i),i} - y_{\pi(i),j}.$$

We know from the theory of difference constraints (see [7]) that such a set of constraints has a solution iff the associated constraint graph G_π has no negative cycles. The constraint graph G_π is a complete weighted directed graph on k vertices $V(G_\pi) = \{1, 2, \dots, k\}$ with each edge (i, j) assigned the weight $w(i, j) = y_{\pi(i),i} - y_{\pi(i),j}$. By definition, the weight of a cycle $(i_1, i_2, \dots, i_p, i_1)$ is given by

$$(y_{\pi(i_1),i_1} - y_{\pi(i_1),i_2}) + (y_{\pi(i_2),i_2} - y_{\pi(i_2),i_3}) + \cdots + (y_{\pi(i_p),i_p} - y_{\pi(i_p),i_1}).$$

Collecting terms, this sum is equal to

$$(y_{\pi(i_1),i_1} + y_{\pi(i_2),i_2} + \cdots + y_{\pi(i_p),i_p}) - (y_{\pi(i_1),i_2} + y_{\pi(i_2),i_3} + \cdots + y_{\pi(i_p),i_1}) \ge 0.$$

It follows that this value is nonnegative, since π is a maximum permutation (matching). Hence, we have proved that if π is a maximum weight permutation, then we can find a vector x that is a splitter point of A with π as the associated directional permutation. Note that a solution vector x can be found by applying an algorithm for shortest paths [7].

For the converse, suppose that π is a directional permutation for A, and x is the associated splitter point. Let $\sigma \neq \pi$ be some other permutation. We have the following inequalities that show that the weight of π is at least as large as the weight of σ:

$$0 = \sum x_i - \sum x_i = \sum_i (x_i - x_{(\pi\sigma^{-1})(i)}) \leq \sum y_{\pi(i),i} - y_{\pi(i),(\pi\sigma^{-1})(i)}$$
$$= \sum y_{\pi(i),i} - \sum y_{\sigma(i),i}.$$

Finally, it follows from the argument above that a permutation π is a unique maximum weight permutation iff there is a strict-splitter point. This can be deduced by replacing the inequalities in the difference constraints with strict inequalities. $\square$

4. Applications to graph connectivity and routing. We now show the connection between the k-directional compass as defined in section 2 and the generation of disjoint paths in graphs using only local information. This connection is established by embedding the vertices of the graph to points of a real vector space. We define a directional embedding of a graph as follows.

DEFINITION. *A graph $G = (V, E)$ has a* directional X-embedding *if there is a mapping f of the vertices to points in real space $f : V \to R^{k-1}$, where $|X| = k$, and there is a k-directional compass $\mathcal{C}_k$, such that for each $v \in V - X$, $f(v)$ is a k-directional splitter point of a subset of $f(N(v))$, the embedded neighbor set of v. A directional X-embedding is* nondegenerate *if, for each $v \in V - X$, $f(v)$ is a strict-splitter point of a subset of $f(N(v))$.*

PROPOSITION 4. *Let $G = (V, E)$ be a graph and let $X \subset V$. If G has a nondegenerate directional X-embedding, then G contains $|X| = k$ independent sink trees to X. Furthermore, these sink trees are strongly independent.*

Proof. As noted earlier, the directional transitivity of the compass implies that paths generated in a consistent direction must be (internally) vertex-disjoint from paths generated consistently in a different direction, when the paths originate from the same vertex. Hence, each of the $|X| = k$ directions can be used to define a distinct parent for each vertex $u \in V - X - \{v\}$. Thus, each of the k directions can be associated with a distinct sink tree, and the collection of k sink trees is independent. It is easily verified that these trees are strongly independent, since we can match any set Y of k points to the k sink trees via the directional splitting permutation π_Y. $\square$

We remark that by viewing a graph as a symmetric digraph we can strengthen Proposition 4 so as to produce k pairwise independent dags; each direction determines a dag with a single sink vertex v. A pair of dags on the same vertex set is independent if, given any vertex u, all paths in one dag are internally node-disjoint from all paths in the other dag originating at the same vertex.

A long-standing, open conjecture of Frank (see [13]) states that all k-connected graphs have, for each vertex v, k-independent spanning trees rooted at v. Our results lead us to conjecture a stronger result.

CONJECTURE 1. *A graph G is k-connected iff for every $X \subset V$, with $|X| = k$, G has a nondegenerate directional X-embedding.*

A closely related conjecture builds on the characterization of [12].

CONJECTURE 2. *A graph G has a convex X-embedding in general position iff G has a nondegenerate directional X-embedding.*

From the results in [12] it follows that the if-part of Conjecture 2 holds. Further, from *st*-numbering we have that the conjecture is true for $|X| = 2$. We prove in section 5 that both conjectures are true for $|X| = 3$. We remark that Conjectures 1 and 2 are not true for directed graphs, since there are k-connected digraphs that do not have k-independent sink trees, for each $k \geq 3$ [9]. However, we now show the conjecture is true for dags k-connected to a sink set. As a corollary of this result we have that dags k-connected to a sink vertex v can be decomposed into k-independent dags (independently proven by [1] and [9]).

THEOREM 2. *A dag k-connected to a set X of k sink vertices has a nondegenerate directional X-embedding.*

Proof. Embed the sink vertices X anywhere in general position in R^{k-1}. It follows from k-connectivity that all nonsink vertices have outdegree at least k. Now topologically sort the remaining vertices. Each vertex, considered in (reverse) topological order, can be positioned as a splitter point of any size k subset of its out-neighborhood. By employing slight adjustments, each point can be made a strict-splitter. Hence, the theorem follows. ☐

5. Directional embeddings of 3-connected graphs. In this section we show that 3-connected graphs have 3-directional embeddings in the plane. To prove this we need a result that is stronger than Theorem 1 (the splitting theorem). Let us assume we are given a 3-directional compass $\mathcal{C}_3(L)$ in the plane equipped with a fixed set of three rays (and the complementary antirays). We now show that it is still possible to "split" a set of three points in the plane into three directions even with the addition of many new directional constraints of a certain type.

LEMMA 1. *Given a pair of points $\{a, c\}$ in the plane, let p be the intersection point of ray-i (of $\mathcal{C}_3(L)$) originating at a with an antiray-j originating at c. Then all the points on the line segment $[a, p]$ are in the same region as a with respect to the reference point c.*

Proof. Clearly, ray-i intersects no other antiray originating at c except for antiray-j. If any point on the segment $[a, p]$ lies in a region different from that of a with respect to the reference point c, then this segment must cross a boundary of directional regions, i.e., it must cross a different antiray from c; this is a contradiction. ☐

LEMMA 2 (the 3-direction constrained splitting lemma). *Let C be any nonempty set of points in the plane. Given a pair of points $\{a, b\}$ in the plane, there is a point $c_0 \in C$ and a point s such that s is a 3-directional splitter point of $\{a, b, c_0\}$ satisfying the following constraints: $d(a, s) = d(a, b)$, $d(b, s) = d(b, a)$, and for each $c \in C$, $d(c, s) = d(c, a)$.*

Proof. Without loss of generality, we can assume that $d(a, b) = 1$ and $d(b, a) = 2$. Consider the pair of rays $\{r, r'\}$, where r is ray-1 originating from a and r' is ray-2 originating at b. Let p denote their point of intersection. Note that the point p is a 3-directional splitter (nonstrict) for $\{a, b, c\}$ (for any $c \in C$) that satisfies the first two directional constraints given in the statement of the lemma, i.e., $d(a, p) = d(a, b)$, $d(b, p) = d(b, a)$. Hence, if no antiray originating at a point of C intersects the ray r, then setting $s = p$ satisfies all the constraints of the lemma. Otherwise, consider all the points of intersection of the ray r with the three antirays originating at points of C. Let p_0 denote the intersection point (of ray r with the antiray originating from a point $c_0 \in C$) that is nearest a. Then by Lemma 1, it follows that $s = p_0$ is a 3-

directional splitter (nonstrict) that satisfies all the constraints and the lemma holds. We remark that by perturbing the points the splitter can be made strict. □

Graphs that are 3-connected have the following characterization well known as Tutte's wheel theorem. A *wheel* is a graph consisting of a cycle (of at least three vertices) and a distinct hub vertex that is connected to every vertex on the cycle.

THEOREM 3 (Tutte's wheel theorem [14]). *If a graph G is 3-connected, then either G is a wheel or there is an edge of G that can be deleted or contracted while preserving 3-connectivity. Furthermore, no contractible edge is part of a triangle.* □

It follows immediately from Tutte's theorem that every 3-connected graph can, through a series of edge contractions and removals of multiedges, be reduced to K_3 while preserving 3-connectivity at each intermediate step. We now extend this result to show that these contractions can be chosen to avoid the edges of any given triangle $S \subset G$.

THEOREM 4 (the reduction theorem). *Let $G = (V, E)$ be a 3-connected graph that contains a triangle S with vertices $\{s_1, s_2, s_3\}$. The graph G can be reduced, via edge contraction and removal of multiedges, to the triangle S, while preserving 3-connectivity at each intermediate stage, and without contracting an edge of S.*

Proof. The proof will follow by induction on the number of edge contractions.

First, suppose that a vertex s_i of the triangle S has only a single neighbor x outside S. The edge $e = s_i x$ can be contracted without affecting 3-connectivity. The 3-connectivity is preserved since the three disjoint paths that all vertices in $V - S$ have to S before the contraction of e remain unaffected after the contraction.

Hence, we can assume all vertices of the triangle S have at least two neighbors in the set $V - S$. Starting with the graph G, consider the sequence of Tutte contractions $\{e_1, e_2, \ldots\}$, and let e_k denote the first edge in the sequence that is not part of the triangle S.

If the graph G_{e_k} obtained by contracting the edge e_k in the graph G is 3-connected, then the theorem follows by induction.

Suppose, on the other hand, that G_{e_k} has a 2-vertex cut C. The 2-cut C must contain a vertex of S, since otherwise C would be a 2-cut in the graph $G' = G_{\{e_1, e_2, \ldots, e_k\}}$ obtained by contracting the sequence of edges $\{e_1, e_2, \ldots, e_k\}$; however, this graph G' is 3-connected by Tutte's theorem. The 2-cut C must also contain a vertex of S, since otherwise C would be a 2-cut in the graph G, which is assumed to be 3-connected. Hence, the 2-cut C in the graph G_{e_k} contains precisely one vertex of S, say, s_3, and one vertex of $V - S$, say, x. No connected component of $G_{e_k}(V - C)$ can contain vertices in both S and $V - S$, since otherwise this would imply that the graph G' has a 2-cut. It follows that the remaining vertices $\{s_1, s_2\}$ in $S - C$ must have the vertex x as their only neighbor in the graph G_{e_k}. Hence, the (original, uncontracted) graph G contains two vertices x_1, x_2 in $V - S$ that are neighbors with each of the two vertices s_1, s_2 of $S - C$, forming a $K_{2,2}$ subgraph.

We claim that we can contract any edge e' of this $K_{2,2}$ subgraph of G without affecting 3-connectivity. This claim follows from the fact that the three disjoint paths to S that each vertex possesses in the graph G can be chosen so that one path uses edge $s_1 x_1$ and another uses $s_2 x_2$. Hence, contracting the edge $s_1 x_1$ (or edge $s_2 x_2$) will not change the number of disjoint paths to S, i.e., the graph $G_{e'}$ is 3-connected. The proof thus follows by induction. □

COROLLARY 1. *Let $G = (V, E)$ be a graph 3-connected to three vertices $S = \{s_1, s_2, s_3\} \subset V$. The graph G can be reduced, via edge contraction and removal of multiedges, to the subgraph induced by S, while preserving 3-connectivity to S at each* *intermediate stage, and without contracting an induced edge of S.*

Proof. The proof is an immediate consequence of Theorem 4, since if we complete S to a triangle, the graph becomes 3-connected. ∎

We apply the reduction theorem to show that we can inductively obtain a 3-directional embedding when a graph is 3-connected to a set of vertices S.

THEOREM 5. *Let G be a graph, and suppose S is any set of three vertices. Then G is 3-connected to S iff G has a nondegenerate directional S-embedding.*

Proof. The proof of the if-part is immediate from Proposition 4. The proof of the converse follows by induction on the number of reduction operations defined by Corollary 1 above.

For the inductive hypothesis we assume we have a nondegenerate directional S-embedding of a reduced graph G'. For the base of the induction embed the three vertices of S anywhere in general position in R^2. Clearly, for an edge addition operation the induction is trivially extended. Edge expansions involving a single vertex of S are also trivial. We need only consider the case of an edge expansion not involving S (i.e., a vertex 3-split operation).

We may assume the vertex $a \notin S$ involved in the edge expansion is a 3-directional splitter of its neighbors represented by $B \cup C$, where $|B| \geq 2 \leq |C|$. Further, we can assume w.l.o.g. that a is a 2-directional splitter of B, i.e., there are two points in B that lie in two distinct directions relative to the reference point a. The proof of the theorem follows from the following claim.

CLAIM. *There exists in the plane a pair of points s_b, s_c such that s_b is a 3-directional splitter of $B \cup \{s_c\}$, and s_c is a 3-directional splitter of $C \cup \{s_b\}$. Furthermore, all previously defined directional constraints towards a are preserved by the new points, i.e., for each $b \in B$, $d(b, a) = d(b, s_b)$ and for each $c \in C$, $d(c, a) = d(c, s_c)$.*

Proof of claim. Set $s_b = a$. A solution for s_c can be found by applying Lemma 2 (the 3-direction constrained splitting lemma), where the point s_c plays the role of s in the statement of Lemma 2. The satisfaction of the first constraint of Lemma 2 ensures that s_b (playing the role of a in the lemma) is a 3-directional splitter of $B \cup \{s_c\}$. The claim follows by the satisfaction of the remaining constraints of Lemma 2. ∎

6. Conclusion. In this paper we proposed a mathematical model for network routing based on generating paths in a consistent direction. Our directional model was developed out of an algebraic and geometric framework. We gave a natural model for defining directions in real spaces and showed that such directions possess a fundamental matching property of combinatorial interest. We defined a generalization of *st*-numberings that was based on embedding the vertices in multidimensional space and applying the defined directions for disjoint path generation. Directional embeddings are motivated by the fact that they encode the disjoint path structure and induce strongly independent sink trees with simple local information. We showed that a dag that is k-connected to a set of sinks has a k-directional embedding in $(k-1)$-space with the sink set as the extreme vertices. Finally, we proved that all 3-connected graphs have 3-directional embeddings in the plane. The problem of whether directional embedding exist for k-connected graphs, for $k > 3$, remains open.

REFERENCES

[1] F.S. ANNEXSTEIN, K.A. BERMAN, AND R. SWAMINATHAN, *A multi-tree generating routing scheme using acyclic orientations*, in Computing and Combinatorics, Proceedings of the Third Annual International Conference, Lecture Notes in Comput. Sci. 1279, Springer-Verlag, New York, 1997, pp. 18–22.

[2] F. Bao, Y. Igarashi, and S.R. Ohring, *Reliable broadcasting in product networks*, Discrete Appl. Math., 83 (1998), pp. 3–20.
[3] D. Bertsekas and R. Gallager, *Data Networks*, Prentice-Hall, New York, 1992.
[4] J. Cheriyan and S.N. Maheshwari, *Finding non-separating induced cycles and independent spanning trees in* 3*-connected graphs*, J. Algorithms, 9 (1988), pp. 507–537.
[5] J. Cheriyan and J. Reif, *Directed s-t numberings, rubber bands, and testing digraph k-vertex connectivity*, Combinatorica, 14 (1994), pp. 435–451.
[6] R. Cohen, B.V. Patel, F. Schaffa, and M. Willebeek-LeMair, *The sink tree paradigm: Connectionless traffic support on ATM LAN's*, IEEE/ACM Trans. on Networking, 4 (1996), pp. 363–374.
[7] T. Corman, C. Leiserson, and R. Rivest, *Introduction to Algorithms*, McGraw-Hill, New York, 1994.
[8] S. Even and R.E. Tarjan, *Computing an st-numbering*, Theoret. Comput. Sci., 2 (1976), pp. 339–344.
[9] A. Huck, *Disproof of a conjecture about independent branchings in k-connected directed graphs*, J. Graph Theory, 20 (1995), pp. 235–239.
[10] A. Itai and M. Rodeh, *The multi-tree approach to reliability in distributed networks*, Inform. and Comput., 79 (1988), pp. 43–59.
[11] A. Lempel, S. Even, and I. Cederbaum, *An algorithm for planarity testing graphs*, in Theory of Graphs, P. Rosentiehl, ed., Gordon and Breach, New York, 1966, pp. 215–232.
[12] N. Linial, L. Lovász, and A. Wigderson, *Rubber bands, convex embeddings, and graph connectivity*, Combinatorica, 8 (1988), pp. 91–102.
[13] A. Schrijver, *Fractional packing and covering*, in Packing and Covering in Combinatorics, A. Schrijver, ed., Mathematical Centre Tracts 106, Mathematisch Centrum, Amsterdam, 1979, pp. 201–274.
[14] W.T. Tutte, *A theory of* 3*-connected graphs*, Indag. Math. (N.S.), 23 (1961), pp. 441–455.
[15] A. Zehavi and A. Itai, *Three tree-paths*, J. Graph Theory, 13 (1989), pp. 175–188.