Chapter 40

Directed s-t Numberings, Rubber Bands, and Testing Digraph k-Vertex Connectivity

Joseph Cheriyan* John H. Reif†

Abstract

Let $G = (V, E)$ be a directed graph and n denote $|V|$. We show that G is k-vertex connected iff for every subset X of V with $|X| = k$, there is an embedding of G in the $(k-1)$-dimensional space $\mathbf{R}^{k-1}$, $f : V \rightarrow \mathbf{R}^{k-1}$, such that no hyperplane contains k points of $\{f(v) \mid v \in V\}$, and for each $v \in V - X$, $f(v)$ is in the convex hull of $\{f(w) \mid (v, w) \in E\}$. This result generalizes to directed graphs the notion of convex embeddings of undirected graphs introduced by Linial, Lovász and Wigderson in "Rubber bands, convex embeddings and graph connectivity," *Combinatorica* 8 (1988), 91–102.

Using this characterization, a directed graph can be tested for k-vertex connectivity by a Monte Carlo algorithm in time $O((M(n) + nM(k)) \cdot (\log n))$ with error probability $< 1/n$, and by a Las Vegas algorithm in expected time $O((M(n)+nM(k)) \cdot k)$, where $M(n)$ denotes the number of arithmetic steps for multiplying two $n \times n$ matrices ($M(n) = O(n^{2.3755})$). Our Monte Carlo algorithm improves on the best previous deterministic and randomized time complexities for $k > n^{0.19}$; e.g., for $k = \sqrt{n}$, the factor of improvement is $> n^{0.62}$. Both algorithms have processor efficient parallel versions that run in $O((\log n)^2)$ time on the EREW PRAM model of computation, using a number of processors equal to $(\log n)$ times the respective sequential time complexities. Our Monte Carlo parallel algorithm improves on the number of processors used by the best previous (Monte Carlo) parallel algorithm by a factor of at least $(n^2/(\log n)^3)$ while having the same running time.

Generalizing the notion of s-t numberings, we give a combinatorial construction of a directed s-t numbering for any 2-vertex connected directed graph.

1 Introduction

The connectivity of directed and undirected graphs is one of the most well studied areas of graph theory. This topic is central to graph theory [B 78], and has diverse applications in computer science, electrical engineering and operations research [E 79], [M 84], [IR 88]. A directed (undirected) graph is said to be k-vertex connected, if it has at least $k+1$ vertices and the deletion of any set of $k-1$ or fewer vertices leaves the graph strongly connected (connected).

The problem of testing a given graph for k-vertex connectivity, because of its importance, has attracted much algorithmic research. Let n and m denote the number of vertices and the number of edges of the graph. For digraphs (directed graphs), the fastest deterministic algorithm known runs in time $O(\max\{k, \sqrt{n}\}k\sqrt{n}m)$ [G 80]; for $k > \sqrt{n}$, the probabilistic technique of [B 82] improves the running time to $O((\frac{\log 1/p}{\log n/k})n^{1.5}m)$, where p denotes the error probability. For undirected graphs, the fastest (deterministic) algorithm runs in time $O(\max\{k, \sqrt{n}\}k^2 n^{1.5})$ [NI 90] and [CKT 91]. Let $M(n)$ denote the number of arithmetic steps for multiplying two $n \times n$ matrices; at present, the best value of $M(n)$ known is $O(n^{2.3755})$ [CW 87]. Recently, [LLW 88] gave a randomized (Monte Carlo) algorithm for testing the k-vertex connectivity of undirected graphs that runs in time $O((M(n) + nM(k)) \cdot (\log n))$ with error probability $< 1/n$. This is faster than the deterministic algorithm for $k \geq n^{0.19}$.

The key idea behind the algorithm of [LLW 88] is a new characterization of k-vertex connected undirected graphs, namely, an undirected graph is k-vertex connected iff for every size-k subset X of V, the graph has a so-called *convex X-embedding* in general position in $\mathbf{R}^{k-1}$, the $(k-1)$-dimensional Euclidean space. See also [LSS 89].

We generalize the notion of convex embeddings to directed graphs, and show that a directed graph is k-vertex connected iff for every size-k subset X of V, the graph has a convex directed X-embedding (defined in Section 2) in general position in $\mathbf{R}^{k-1}$.

There is a key difference between our proof and the approach of Linial, Lovász and Wigderson [LLW 88].

*School of O. R. and I. E., Cornell University, Ithaca, NY 14853–3801. Supported in part by the Lucille and David Packard Fellowship of Éva Tardos.

†Department of Computer Science, Duke University, Durham, NC 27706. Supported in part by Air Force Contract AFOSR-87-0386, DARPA/ISTO contract N00014-88-K-0458 and N00014-91-J-1985, NASA subcontract 550-63 of primecontract NAS5-30428.

They use a potential or "quadratic energy function", i.e., they associate an energy with each embedding of G, and then use the strict convexity of the potential to argue that the minimum potential determines a unique embedding. The minimum embedding is computed by equating the gradient of the potential to zero, and solving the resulting system of linear equations. Unfortunately, this approach does not seem to work for directed graphs because the potential is independent of *arc directions*. Consequently, the minimum embedding is not dependent on the directions of the arcs; in contrast, for a convex directed embedding only the outgoing arcs of a vertex determine its position, and not the incoming arcs. Instead of working with the potential, we work directly with the system of linear equations that determines the embedding, and we ensure that the embedding is unique by showing that the matrix associated with the system of equations is nonsingular.

A convex directed embedding can be computed with high probability by assigning random weights to the edges and solving a system of linear equations whose coefficients are determined by the random weights. Based on this method, and using the approach of [LLW 88], we develop efficient randomized algorithms for testing a directed graph for k-vertex connectivity. Specifically, we have a Monte Carlo algorithm that runs in time $O((M(n) + nM(k)) \cdot (\log n))$ with error probability $< 1/n$. This algorithm may err both ways: It may reject a k-vertex connected directed graph and may accept a directed graph that is not k-vertex connected. We also have a Las Vegas algorithm that runs in expected time $O((M(n) + nM(k)) \cdot k)$; this algorithm never errs. Our Monte Carlo algorithm improves on the best deterministic algorithm known for all edge densities for $k > n^{0.19}$ by a factor of at least $k^2/n^{0.38}$; e.g., for $k = \sqrt{n}$, the factor of improvement is $> n^{0.62}$. Even for smaller k, $k = 2, 3, \ldots$, our algorithm is faster for relatively dense directed graphs (i.e., for $m > n^{1.38}/k$).

Both our algorithms have efficient parallel implementations, and our Monte Carlo algorithm substantially improves on the best previous parallel algorithm. On the Exclusive Read Exclusive Write (EREW) PRAM model of parallel computation, both algorithms run in $O((\log n)^2)$ time. The Monte Carlo algorithm uses $((M(n)+nM(k)) \cdot (\log n)^2)$ processors, and the Las Vegas algorithm uses $((M(n)+nM(k)) \cdot k \log n)$ processors. For general k, the best previous parallel algorithm for testing the k-vertex connectivity of directed graphs is a Monte Carlo algorithm that runs in $O((\log n)^2)$ time and uses at least $nP(n,m)$ processors. Here, $P(n,m)$ is the number of processors needed to find a maximum matching on an n-vertex m-edge undirected graph in $O((\log n)^2)$ time with high probability. Currently, the best value known for $P(n,m)$ is $O(nmM(n))$ [MVV 87]. For $k = 1$, the best deterministic parallel algorithm known runs in $O((\log n)^2)$ time and uses $M(n)$ processors. For general k, our Monte Carlo algorithm improves on the number of processors used by the best previous algorithm by a factor of at least $(nm/(\log n)^3)$ while having the same running time. Unless k is relatively high, the factor of improvement is even greater : e.g., for $k = \sqrt{n}$, the factor of improvement is $(n^2m/(logn)^3)$.

For undirected graphs, a non-degenerate convex embedding for $k = 2$ is equivalent to a so-called s-t numbering of the vertices. Lempel, Even and Cederbaum introduced the notion of s-t numberings of 2-vertex connected undirected graphs [LEC 66]. Since then, this notion has been used for planarity testing [LEC 66], fault-tolerant protocols for distributed computers [IR 88], etc. It appears that the obvious generalization to directed graphs has not been studied before. We give a simple combinatorial proof showing that every 2-vertex connected directed graph has a directed s-t numbering (defined in Section 3). This result is implied by our main theorem that shows that every k-vertex connected directed graph has a non-degenerate convex directed embedding, however, the result for $k = 2$ may be of independent interest since our proof uses straightforward combinatorial techniques. For instance, the proof for $k = 2$ gives a deterministic algorithm, whereas the algorithm given by the main proof uses randomization.

An immediate application of directed s-t numberings is to give a simple construction for two independent branchings (defined in Section 2) of any 2-vertex connected directed graph. Previously, a much more complicated construction has been reported by Whitty [W 87]; very recently, we have learned of another construction due to Plehn [P 91].

In Section 3, we focus on 2-vertex connected directed graphs, and give a combinatorial construction for directed s-t numberings as well as the construction for two independent branchings, and in Section 4 we give an algebraic proof of our main result. Section 5 gives our randomized algorithms for testing a directed graph for k-vertex connectivity.

2 Preliminaries

This section contains only definitions and notation.

Let $G = (V, E)$ be a digraph. For any $v \in V$, call the set $\{w : (v, w) \in E\}$ the *successors* of v, and denote this set by $\Gamma(v)$. Call the set $\{u : (u, v) \in E\}$ the *predecessors* of v. Define the *reversal* of G, rev(G), to be the digraph $(V, \{(w, v) : (v, w) \in E\})$.

For any $X \subset V$ with $|X| = k$, define a *convex directed X-embedding* of G to be a function $f : V \rightarrow \mathbf{R}^{k-1}$, such that for each $v \in V - X$, $f(v)$ is in the convex hull of $\{f(w) \mid (v, w) \in E\}$. f is an embedding of the (vertices of) digraph G in the $k-1$ dimensional space $\mathbf{R}^{k-1}$.

A *branching* B is a spanning acyclic subgraph of G such that one vertex, called the *root*, has outdegree $= 0$ and all other vertices have outdegree $= 1$. Note that every vertex has exactly one directed path to the root in B. For any fixed vertex z of a digraph, two branchings rooted at z are called *independent* if they are edge disjoint, and further, for each vertex v, the two paths from v to z in the two branchings are openly vertex disjoint. For $k > 2$, k branchings are called *independent* if they are pairwise independent.

For subsets X and Y of V, $p(X, Y)$ denotes the maximum number of vertex disjoint paths from X to Y. A *separator* is a subset S of V such that $G - S$ is not strongly connected. For vertices v and z, a *v-z separator* is a subset S of $V - \{v, z\}$ such that $G - S$ has no path from v to z.

For a set of points X in $\mathbf{R}^k$, $\mathrm{conv}(X)$ denotes the convex hull of X, and $\mathrm{rank}(X)$ denotes the affine rank of X, i.e., one plus the dimension of the smallest affine space containing X.

Our complexity bounds are stated in the uniform-cost model, i.e., an arithmetic operation (addition, subtraction, multiplication or division) or a comparison on the integers or on the integers modulo a prime counts as a single step. The largest integer computed by our algorithms has value $n^{O(1)}$, and can be represented using $O(\log n)$ bits. To obtain the bounds in the logarithmic-cost or bit-complexity model, multiply the sequential running times and the number of parallel processors by a factor of $(\log n)^2$ (this factor is not optimal, but suffices for us).

For our probability estimates below, we assume that n is not too small, e.g., $n \geq 64$ would suffice.

3 *s-t* numberings for directed graphs

We generalize the notion of *s-t* numberings to 2-vertex connected digraphs, and show the following.

THEOREM 3.1. *Let $G = (V, E)$ be a 2-vertex connected directed graph, and s and t be any two vertices of G. Then, there is a numbering $\pi : V \rightarrow \{1, \ldots, n\}$ of V (i.e., π is a bijection) such that*
(i) $\pi(s) = 1$ and $\pi(t) = n$, and
(ii) for each $v \in V - \{s, t\}$, there is a successor w with $\pi(w) > \pi(v)$ and a successor u with $\pi(u) < \pi(v)$.

The numbering π is called a *directed s-t numbering*.

Instead of proving Theorem 3.1 directly, we prove a generalization.

THEOREM 3.2. *Let $G = (V, E)$ be a directed graph, and z be a vertex of G such that for each $v \in V - \{z\}$, G has two openly vertex disjoint paths from v to z. Let Z denote the set $\{z\} \cup \{y : (y, z) \in E\}$, i.e., the set containing z and its predecessors. Then, there is a numbering $\pi : V \rightarrow \{1, \ldots, n\}$ of V such that for each $v \in V - Z$, there is a successor w with $\pi(w) > \pi(v)$ and a successor u with $\pi(u) < \pi(v)$.*

Note that the highest numbered vertex and the lowest numbered vertex according to π must belong to Z.

To prove Theorem 3.1, add a new vertex z and the edges (s, z) and (t, z) to the original 2-vertex connected digraph, and invoke Theorem 3.2 on the augmented digraph.

Proof. (of Theorem 3.2) The proof uses induction on n. The basis of the induction is easily seen to hold, so we proceed with the induction step.

Using a theorem due to Lovász [Lo 73], it can be shown that G has a spanning subgraph $G' = (V, E')$ such that the outdegree of z (in G') is $= 0$, and for each $v \in V - \{z\}$, the outdegree of v (in G') is $= 2$, and G' has two openly vertex disjoint paths from v to z. Clearly, any valid numbering π of G' is also a valid numbering of G. Focus on the in-degree of the vertices in G'. Since, G' has at most $2n - 4$ edges with both end vertices in $V - \{z\}$ (there must be ≥ 2 edges entering z), G' has a vertex with in-degree ≤ 1.

If G' has a vertex v with in-degree $= 0$, then it can be seen that $G' - v$ satisfies the condition of the theorem, therefore we can find a numbering π' for $G' - v$ and then extend it to a numbering π for G (e.g., by assigning v a number between those of its successors, and then "shifting up" some numbers by one in order to avoid duplication of $\pi(v)$).

If G' has a vertex v with exactly one entering edge, say (u, v), then we contract this edge (i.e., replace vertices u and v by the vertex v^*, and replace edges (v, w) or (u, w) by (v^*, w) and replace edges (q, u) by (q, v^*)) and consider the resulting digraph $G'/(u, v)$. By applying Menger's theorem, it can be seen that $G'/(u, v)$ satisfies the condition of the theorem, therefore we can find a numbering π' for $G'/(u, v)$ and then extend it to a numbering π for G (e.g., by assigning $\pi(u) := \pi'(v^*)$, and assigning v a number between the numbers of its two successors which is either less than or greater than $\pi(u)$, depending on $\pi'(w)$, where w is the other successor of u, and then "shifting up" or "shifting down" some numbers by one in order to avoid duplication of $\pi(v)$).

One application of Theorem 3.1 is to construct for any 2-vertex connected digraph with a specified root vertex $z \in V$, two independent branchings rooted at z, i.e., to find branchings $B_1 = (V, E_1)$ and $B_2 = (V, E_2)$ rooted at z such that $E_1 \cap E_2 = \emptyset$ and for each $v \in V$ the two paths from v to z, respectively in B_1 and B_2, are openly vertex disjoint. For $k = 2$, this gives a simple solution for the independent branchings conjecture: For any vertex z of a k-vertex connected directed graph, there are k independent branchings rooted at z. The conjecture was first raised by A. Frank, see [S 79, pp. 235], and later Itai and Rodeh [IR 88] raised a variant of the conjecture for undirected graphs. For $k \leq 3$, the conjecture for undirected graphs has been resolved by [CM 88] and [ZI 89]. For general k, the conjecture remains open for directed graphs for $k > 2$ and for undirected graphs for $k > 3$.

A construction for two independent branchings has been presented before in [W 87]. Our construction, which is much simpler than the complicated one in [W 87], is straightforward:
First, compute an s-t numbering for G with $s = z$ and $t = y$, where y is any predecessor of z. For the "low" branching B_1, each $v \in V - \{z\}$ chooses its parent to be a lower numbered successor, and for the "high" branching B_2, each $v \in V - \{z, y\}$ chooses its parent to be a higher numbered successor (B_2 also contains the edge (y, z)).

THEOREM 3.3. (WHITTY) *A digraph G is 2-vertex connected iff for every vertex z, G has two independent branchings rooted at z.*

4 Convex directed embeddings

Our main result here is that a digraph is k-vertex connected iff for every subset X of V, with $|X| = k$, G has a non-degenerate convex directed X-embedding, i.e., the vertices can be embedded in $\mathbf{R}^{k-1}$ such that they are in general position and each vertex $v \in V - X$ is in the convex hull of the successors of v.

Intuitively, it is easy to construct a convex directed X-embedding of G in $\mathbf{R}^{k-1}$, where X is any size-k subset of V : Embed the vertices of X at the corners of the unit simplex, let each arc behave as a "directed rubber band", and compute the equilibrium positions of the vertices. In more detail, let each arc exert zero force on its head vertex, and a force on its tail vertex equal to the length of the arc times an arbitrary positive coefficient. Note that the asymmetry in the behavior of "directed rubber bands" is critical for obtaining a convex directed embedding. This physical system of directed rubber bands can be easily modeled as a system of linear equations, and the equilibrium position of each vertex $v \in V - X$ can be found by solving the system.

Let $c : E \to \mathbf{R}$ be a positive real-valued coefficient function for the edges, i.e., the "directed rubber band" (u, v) has coefficient of elasticity $c(u, v)$. Let $v_1, \ldots, v_{n-k}$ and $v_{n-k+1}, \ldots, v_n$ be arbitrary linear orderings of $V - X$ and X, respectively. Using c, we can find an embedding $f^{(c)} : V \to \mathbf{R}^{k-1}$ with the vertices of X at the corners of the unit simplex, i.e., with $f^{(c)}(v_n) = (0, \ldots, 0)$ and for $j = 1, \ldots, k-1$, $f^{(c)}(v_{n-k+j}) = e_j$, where e_j is the vector in $\mathbf{R}^{k-1}$ with a "1" in position j and zeros in all other positions. Moreover, $f^{(c)}$ satisfies the following equations, where we denote the coordinates of $f^{(c)}(v)$ by $f_1^{(c)}(v), f_2^{(c)}(v), \ldots, f_{k-1}^{(c)}(v)$. For notational convenience, we use f instead of $f^{(c)}$, when there is no danger of confusion.
For all $v \in V - X$ and all $j = 1, \ldots, k-1$,

$$\sum_{w \in \Gamma(v)} c(v, w) \cdot (f_j(w) - f_j(v)) = 0. \qquad (1)$$

Intuitively, this embedding f ensures that each vertex in $V - X$ is in equilibrium under the action of the forces exerted by the directed rubber bands.

The system (1) can be rewritten as

$$AF_j = B_j \qquad j = 1, \ldots, k-1, \qquad (\star)$$

where A is the $(n-k) \times (n-k)$ matrix whose ijth entry, $i, j = 1, \ldots, n-k$, is given by

$$\begin{aligned} A_{ij} &= -c(v_i, v_j) && \text{if } i \neq j \text{ and } (v_i, v_j) \in E, \\ &= \sum_{v_l \in \Gamma(v_i)} c(v_i, v_l) && \text{if } i = j, \\ &= 0 && \text{otherwise;} \end{aligned}$$

B_j and F_j are column vectors of length $n - k$ with the ith entry ($1 \leq i \leq n-k$) of F_j, F_{ij}, standing for the jth co-ordinate of v_i, $f_j(v_i)$, and with

$$\begin{aligned} B_{ij} &= c(v_i, v_{n-k+j}) && \text{if } (v_i, v_{n-k+j}) \in E, \\ &= 0 && \text{otherwise.} \end{aligned}$$

Clearly, f is a convex directed X-embedding, i.e., for each vertex $v \in V - X$, $f(v)$ is in the convex hull of $f(\Gamma(v)) = \{f(w) \mid w \in \Gamma(v)\}$, otherwise there would be a nonzero force $\sum_{w \in \Gamma(v)} c(v, w) \cdot (f(w) - f(v))$ acting on v.

The next lemma shows that the matrix A is nonsingular, consequently equation $(\star)$ determines a unique embedding f.

Our main result, Theorem 4.2, states something more: For the embedding f and every subset U of V, the (affine) dimension of $f(U) = \{f(w) \mid w \in U\}$ plus one is exactly equal to the number of vertex disjoint paths between U and X, $p(U, X)$. First, we show in Theorem 4.1 that for any convex directed X-embedding f, the dimension of $f(U)$ plus one is less than or equal to $p(U, X)$.

LEMMA 4.1. *If G is k-vertex connected, then the matrix A is nonsingular.*

Proof. The matrix A has the following key property:
For each i, $i = 1, \ldots, n-k$,

$$|A_{ii}| \geq \sum_{\substack{1 \leq j \leq n-k \\ j \neq i}} |A_{ij}|. \tag{2}$$

Further, there is at least one row for which the inequality is strict, because any vertex $v_{n-k+q} \in X$ $(1 \leq q \leq k)$ has at least k predecessors so it has a predecessor, say, $v_h \in V - X$, and therefore

$$|A_{hh}| = \sum_{(v_h, w) \in E} c(v_h, w) > \sum_{\substack{1 \leq j \leq n-k \\ j \neq h}} |A_{hj}|. \tag{3}$$

To obtain a contradiction, suppose that A is singular. Then, clearly there exists $x \in \mathbf{R}^{n-k}$ such that

$$Ax = 0,$$

otherwise, the nullity of A would be zero and A would have full rank.
By considering the inner product of the hth row of A, denoted a_h, with x,

$$a_h \cdot x = 0,$$

it can be seen that all entries of x do not have the same absolute value, otherwise

$$\begin{aligned} |A_{hh}||x_h| &> |x_h| \cdot \sum_{\substack{1 \leq j \leq n-k \\ j \neq h}} |A_{hj}| && \text{using (3)} \\ &\geq \Big| \sum_{\substack{1 \leq j \leq n-k \\ j \neq h}} A_{hj} x_j \Big| \\ &= |A_{hh}||x_h|. \end{aligned}$$

In other words, $|x_h| < |x_j|$, for some $j \in \{1, \ldots, n-k\}$. After renumbering the vertices if necessary, it follows that there is an l $(l < n-k)$ such that:

$$|x_1| = \cdots = |x_l| > |x_{l+1}| \geq \cdots \geq |x_{n-k}|. \tag{4}$$

Focus on the first l rows of A, and the corresponding vertices $v_1, \ldots, v_l$ of G. Since G is k-vertex connected, at least one of the vertices $v_1, \ldots, v_l$ has a successor in $V - \{v_1, \ldots, v_l\}$. Consequently, at least one of these rows, say, the qth row has

$$|A_{qq}| > \sum_{\substack{1 \leq j \leq l \\ j \neq q}} |A_{qj}|. \tag{5}$$

Now, using the inner product of the qth row of A with x, $a_q \cdot x = 0$, we obtain the desired contradiction because

$$\begin{aligned} |A_{qq}||x_q| &\geq |x_q| \cdot \sum_{\substack{1 \leq j \leq n-k \\ j \neq q}} |A_{qj}| && \text{using (2)} \\ &> \sum_{\substack{1 \leq j \leq n-k \\ j \neq q}} |A_{qj}||x_j| && \text{using (4) and (5)} \\ &\geq \Big| \sum_{\substack{1 \leq j \leq n-k \\ j \neq q}} A_{qj} x_j \Big| \\ &= |A_{qq}||x_q|. \end{aligned}$$

REMARK. Notice that the hypothesis of the lemma can be weakened without affecting the proof: To ensure that A is nonsingular, G should be strongly connected, and each vertex of X should have at least k predecessors.

THEOREM 4.1. *Let $G = (V, E)$ be a digraph, and X be a subset of V. Then for every convex directed X-embedding $f : V \to \mathbf{R}^{k-1}$ of G and every subset U of V, $U \neq \emptyset$,* $\operatorname{rank}(f(U)) \leq p(U, X)$.

Proof. Consider any fixed subset U of V. Menger's theorem implies that there is a subset S of V with $|S| = p(U, X)$ such that $G - S$ has no path from any vertex of U to any vertex of X. Let W denote the maximum subset of $V - S$ such that $G - S$ contains no path from any vertex of W to any vertex of X. For any convex directed X-embedding f, the following argument shows that each extreme point of $\operatorname{conv}(f(W \cup S))$ belongs to $f(S)$. Consider any vertex u in W, and note that by the definition of W, $u \notin X$. Since f is a convex directed embedding it follows that $f(u)$ is contained in $\operatorname{conv}(f(\Gamma(u)))$, consequently, since $\Gamma(u)$ is a subset of $W \cup S - \{u\}$, $f(u)$ is contained in $\operatorname{conv}(f(W \cup S - \{u\}))$. This shows that $\operatorname{conv}(f(W \cup S)) = \operatorname{conv}(f(S))$. It now follows that $\operatorname{rank}(f(U)) \leq \operatorname{rank}(f(W \cup S)) = \operatorname{rank}(f(S)) \leq |S| = p(U, X)$.

THEOREM 4.2. *Let $G = (V, E)$ be a digraph, and X be a subset of V. Then, G has a convex directed X-embedding $f : V \to \mathbf{R}^{k-1}$ such that for every subset U of V, $U \neq \emptyset$,* $\operatorname{rank}(f(U)) \geq p(U, X)$.

Proof. Consider a fixed but arbitrary subset U' of V, and focus on $p(U', X)$. If $p(U', X) \geq k$, then let $U = \{u_1, u_2, \ldots, u_k\}$ denote any subset of U' such that G has k vertex disjoint paths from U to X, i.e., the ith path starts with u_i and ends with a distinct vertex of X. Otherwise, if $p(U', X) = q < k$, then let $U = \{u_1, \ldots, u_q\} \cup (X - \{x_1, \ldots, x_q\})$ denote any subset of $U' \cup X$ obtained by taking the set of start vertices $\{u_1, \ldots, u_q\} \subset U'$ of q vertex disjoint paths from U' to X, and adding all the vertices of X *except* the vertices $x_1, \ldots, x_q$ that belong to these q vertex disjoint paths.

Let the vertices of U be denoted by $v_1, \ldots, v_k$. Observe that $p(U, X) = k$, by the construction of U.
Let $c : E \to \mathbf{R}$ be a vector of positive coefficients, and $f^{(c)}$ be the embedding obtained by solving equation ($\star$). The rank of $f^{(c)}(U)$ is determined by the matrix $g(c, U)$ given by

$$g(c,U) = \begin{pmatrix} 1 & f_1^{(c)}(v_1) & \ldots & f_{k-1}^{(c)}(v_1) \\ 1 & f_1^{(c)}(v_2) & \ldots & f_{k-1}^{(c)}(v_2) \\ \vdots & \vdots & & \vdots \\ 1 & f_1^{(c)}(v_k) & \ldots & f_{k-1}^{(c)}(v_k) \end{pmatrix}.$$

The rank of $f^{(c)}(U)$ is $< k$ exactly when the matrix $g(c, U)$ is singular, i.e., exactly when the determinant of $g(c, U)$ is zero. Note that $\det(g(c, U))$ is a rational function over the entries $c_1, \ldots, c_m$ of c, since each $f_j^{(c)}(u_i)$ is an element of the matrix $A^{-1}B_j$ (see equation ($\star$)). Therefore, either the determinant is identically zero (for all c), or it becomes zero only for a set of vectors c of measure zero.
We claim that the first possibility is ruled out, i.e., $\det(g(c, U))$ is not identically zero. The following argument shows that the coefficients c may be chosen such that for the resulting embedding $f^{(c')}$ the affine dimension of $f^{(c')}(U)$ plus one equals $p(U, X)$: Let $P_1, \ldots, P_k$ be vertex disjoint paths from U to X; suppose that the path P_i has start vertex u_i and end vertex x_i ($1 \leq i \leq k$). Choose $c'(u, v) = \gamma$ for each arc (u, v) in each of the paths $P_1, \ldots, P_k$, and choose $c'(u, v) = 1$ for the remaining arcs. We claim that for any given $\epsilon > 0$ there exists a sufficiently large γ such that for each path P_i the distance of u_i from x_i is less than ϵ. The claim would imply that the dimension of $f^{(c')}(U)$ plus one equals k, because the k vertices of U (that are the start vertices of the paths $P_1, \ldots, P_k$) would be located arbitrarily near to k distinct vertices of the unit simplex. To prove the claim, focus on a fixed P_i and let the vertex sequence of P_i be $(y_1 = u_i), y_2, \ldots, y_{l-1}, (y_l = x_i)$. Suppose that the equilibrium positions of y_{l-1} and $y_l = x_i$ are given by $f^{(c')}(y_{l-1})$ and $f^{(c')}(y_l)$. Then, the vector $\gamma \cdot (f^{(c')}(y_l) - f^{(c')}(y_{l-1}))$ gives the magnitude and the direction of the force exerted on y_{l-1} by the arc (y_{l-1}, y_l). The remaining force on y_{l-1} has a total magnitude of at most $\sqrt{2}n$, since there are at most n remaining arcs with tail vertex y_{l-1}, and each of these arcs has a length of at most $\sqrt{2}$ and a coefficient of 1. As γ increases, the equilibrium distance of y_{l-1} and x decreases, and for a sufficiently high γ the distance becomes less than ϵ/n. Repeating the above argument for the other vertices $y_{l-2}, \ldots, y_2, (y_1 = u_i)$ of P_i, it follows that all the vertices, including u_i, are within a distance ϵ of x_i.
Also, by the construction of U from U', if $f^{(c')}(U)$ has rank $= p(U, X) = k$, then $f^{(c')}(U')$ has rank $\geq p(U', X)$. Therefore, for each subset U' of V, the measure of the set of vectors c such that $\text{rank}(f^{(c)}(U'))$ is less than $p(U', X)$ is zero. The theorem now follows, because there are at most 2^n sets U' (i.e., finitely many sets), therefore there exists a c with $\text{rank}(f^{(c)}(U')) \geq p(U', X)$ for all subsets U' of V.

THEOREM 4.3. *Let $G = (V, E)$ be a digraph, and k be an integer, $k = 2, \ldots, n-1$. Then, G is k-vertex connected iff for every subset X of V with $|X| = k$, G has a convex directed X-embedding in general position.*

Proof. Let X be any subset of V with $|X| = k$. Since G is k-vertex connected, it is clear that for every subset Y of V with $|Y| = k$, $p(Y, X) = k$. Now, the above theorem guarantees a convex directed X-embedding $f : V \to \mathbf{R}^{k-1}$ such that for every subset Y of V with $|Y| = k$ $\text{rank}(f(Y)) \geq p(Y, X) = k$. In other words, no hyperplane contains k vertices of $f(V)$, therefore, f is in general position.
For the other direction of the theorem, let f be a convex directed X-embedding of G in general position, where X is any subset of V with $|X| = k$. Then, for every $Y \subset V$ with $|Y| = k$, $p(Y, X)$ is $\geq \text{rank}(f(Y))$ by Theorem 4.1, and since f is in general position $\text{rank}(f(Y)) = k$. Consequently, G is k-vertex connected.

To gain computational efficiency, instead of finding an X-embedding f in $\mathbf{R}^{k-1}$, we use the [LLW 88] method of computing an X-embedding f in the $k-1$ dimensional linear space over a finite field F; the resulting f is called a *modular* X-embedding. The computation is done over the field of integers modulo a prime p, Z_p. A *random modular* X-embedding is constructed as follows : Fix a random prime p in the range $[n^5, n^6]$ and do the computations over $(Z_p)^{k-1}$. Choose a random nonzero coefficient function $c : E \to Z_p$ on the edges, i.e., the coefficients are drawn from the nonzero elements of Z_p, and compute an X-embedding $f : V \to (Z_p)^{k-1}$ by solving equation ($\star$). Two remarks on computing random modular X-embeddings are in order : Firstly, the matrix A in equation ($\star$) may be singular over Z_p, i.e., possibly $\det(A) = 0 \bmod p$, even though A is nonsingular over the field of integers. Since p is drawn randomly from the primes in the range $[n^5, n^6]$, and $\det(A) \neq 0$ over the integers, the probability of A being singular over Z_p is less than $1/n^4$. * Secondly, for a subset U of $V - X$, there is a positive probability that a randomly chosen nonzero coefficient function c on the edges gives a modular X-embedding $f^{(c)}$ with $\text{rank}(f^{(c)}(U)) < p(U, X)$. Consequently, for a k-vertex connected digraph G, it is *not* necessarily true that there exists a modular X-embedding f such that *every* subset

*The number of distinct prime divisors of $\det(A)$ is $\leq \log|\det(A)| \leq \log(np)^n \leq 7n\log n$.

U of V with cardinality k has $\text{rank}(f(U)) = k$.

LEMMA 4.2. *If G is k-vertex connected then for any fixed subset U of $V-X$ with $|U| = k$, a random modular X-embedding $f : V \to (Z_p)^{k-1}$ satisfies* $\text{rank}(f(U)) = k$ *with probability at least* $1 - (1/n^3)$.

Proof. Consider a vector c of nonzero coefficients drawn at random from $(Z_p - \{0\})^{|E|}$. With probability $> 1 - \frac{1}{n^4}$, the matrix A in equation $(\star)$ has $\det(A) \neq 0 \bmod p$. Moreover, assuming that $\det(A) \neq 0 \bmod p$, the determinant $\det(g(c,U))$ in the proof of Theorem 4.2 is a rational function of degree $\leq \frac{3}{4}n^2$ whose denominator is nonzero, because each element of the matrix $g(c,U)$ is a rational function of degree $\leq 2(n-k)$ whose denominator (namely, $\det(A)$) is nonzero. Note that each element is obtained by solving equation $(\star)$. Now applying the Zippel-Schwartz lemma [Z 79], [Sc 80], the probability that $\det(g(c,U)) = 0 \bmod p$ (assuming that $\det(A) \neq 0 \bmod p$) is at most $\frac{3n^2}{4p}$. Therefore, the probability that $\text{rank}(f^{(c)}(U))$ is less than k, which is at most the probability that either $\det(A) = 0 \bmod p$ or $\det(g(c,U)) = 0 \bmod p$, is at most $\frac{1}{n^4} + \frac{3n^2}{4p} \leq \frac{1}{n^3}$.

5 Algorithms for digraph k-vertex connectivity

Our algorithms are similar to those of [LLW 88]. We first give a Monte Carlo algorithm for testing a digraph for k-vertex connectivity, and then describe the modifications necessary to obtain a Las Vegas algorithm.

The basic subroutine used by our algorithms, *testroot*, has two inputs, a digraph $G = (V,E)$ and a "root vertex" $z \in V$. The subroutine tests whether G has k vertex disjoint paths to z from each $v \in V - X$, where X is a set of k predecessors of z. A call of *test root* is said to succeed if the call finds that G has the above paths, otherwise, the call is said to fail.

For every vertex $v \in V$, let $\Gamma_k(v)$ denote a fixed but arbitrarily chosen set of k successors of v, i.e., $\Gamma_k(v) \subseteq \Gamma(v)$ and $|\Gamma_k(v)| = k$.

The subroutine *testroot*, given a root vertex z, works as follows: First, an arbitrary subset X of the predecessors of z of size k is chosen. Then, a random modular X-embedding $f : V \to (Z_p)^{k-1}$ is computed by choosing a random prime p and random nonzero coefficients for the edges of G (i.e., choosing a random nonzero function $c : E \to Z_p$) and solving the linear system $(\star)$. Finally, for every $v \in V-X$, we test whether $\text{rank}(f(\Gamma_k(v)))$ is equal to k. If some vertex v fails the test, then the call of *testroot* fails, otherwise the call succeeds. See Figure 1.

The following well known result [E 75,E 79] shows how *test root* can be used to test for k-vertex connectivity.

LEMMA 5.1. *Let $z_1,\ldots,z_k$ be k arbitrary but distinct vertices of the digraph G. Then, G is k-vertex connected iff for both the digraphs G and* $\text{rev}(G)$, *testroot succeeds with each of the k root vertices $z = z_1,\ldots,z_k$.*

For each execution of our Monte Carlo algorithm, if we do not require the output to be correct but instead allow a digraph that is *not* k-vertex connected to be accepted with probability $< 1/n$, then we can improve on the number of calls to *testroot* by fixing an appropriate value $k' < k$ and calling *testroot* for k' randomly chosen root vertices, say $y_1,\ldots,y_{k'}$. An execution of the algorithm may be erroneous: If G is not k-vertex connected, but every separator of G with size $< k$ contains the randomly chosen root vertices $y_1,\ldots,y_{k'}$, then the algorithm would accept G. To fix the value of k', note that the probability of an erroneous result is $\leq ((k-1)/n)^{k'}$. This idea has been used before by [B 82], [M 84] and [LLW 88].

LEMMA 5.2. *Let k' be an integer such that $((k-1)/n)^{k'} < 1/n$; note that k' may be taken to be $\lceil \log n/\log(n/k) \rceil$. Let $y_1,\ldots,y_{k'}$ be k' randomly chosen vertices of the digraph G. For both the digraphs G and* $\text{rev}(G)$, *if test root succeeds with each of the k' root vertices $z = y_1,\ldots,y_{k'}$, then with probability $> 1-(1/n)$ G is k-vertex connected.*

As mentioned before, with probability $< 1/n$ the Monte Carlo algorithm may accept a digraph that is not k-vertex connected. Also, with probability $< 1/n$ the algorithm may reject a digraph that is k-vertex connected. This is because with probability $< 1/n$ the embedding f computed by the algorithm may have a vertex v with $\text{rank}(f(\Gamma_k(v))) < k$, even though the digraph is k-connected. See Figure 2.

Coming to the running time analysis, the linear system $(\star)$ has to be solved in order to construct a modular embedding : We first compute the inverse of A in time $O(M(n-k))$, and then compute F_j for $j = 1,\ldots,k$ "simultaneously", by multiplying A^{-1} and the matrix $(B_1 \ldots B_k)$ in time $O(M(n-k)\cdot\lceil k/(n-k)\rceil)$. For subsets U of V with $|U| = k$, the running time for computing $\text{rank}(f(U))$ is $O(M(k))$. Therefore, *testroot* runs in time $O(M(n-k)\lceil k/(n-k)\rceil + M(k)(n-k))$. To estimate the cost of $\log n/\log(n/k)$ calls of *test root*, note that when $k \leq n/2$, then $\log n/\log(n/k) \leq \log n$, otherwise, $\log n/\log(n/k) \leq k \log n/(n-k)$.

The next theorem sums up the above discussion.

THEOREM 5.1. *The k-vertex connectivity of any digraph can be tested by a Monte Carlo algorithm with a running time of $O((M(n)+nM(k))\cdot(\log n))$. If the digraph is k-connected (is not k-connected), then the algorithm accepts the digraph (rejects the digraph) with probability $> 1-1/n$.*

A Las Vegas algorithm for testing the k-vertex

```
Subroutine testroot(z, G)
Input: a digraph G = (V, E), and a "root vertex" z.
Output: for a subset X of k predecessors of z, if there is a v ∈ V − X
          such that G has < k vertex disjoint paths from v to z, then
          return failure, otherwise return success with probability > 1 − 1/n.
begin
    choose a random prime p ∈ [n^5, n^6] and
    a random coefficient vector c ∈ (Z_p − {0})^|E|;
    let X be a set of k predecessors of z;
    compute the modular X-embedding f^(c) by solving equation (⋆);
    for all v ∈ V − X do
        compute the rank of f^(c)(Γ_k(v)) = {f^(c)(w) | w ∈ Γ_k(v)};
    if matrix A is singular over Z_p or some v computes a rank < k
        then return failure
        else return success;
end.
```

Figure 1: Constructing and testing a modular X-embedding

```
Algorithm Monte Carlo k-connectivity
Input: a digraph G = (V, E).
Output: if G is k-connected, accept it with probability > 1 − 1/n,
          otherwise reject it with probability > 1 − 1/n.
begin
    let k' = ⌈log n / log(n/k)⌉;
    for i = 1 to k' do
    begin
        Choose a random vertex y_i ∈ V;
        Call testroot(y_i, G) and testroot(y_i, rev(G)) and
        if either call of testroot fails, then reject G and stop;
    end;
    accept G;
end.
```

Figure 2: Monte Carlo algorithm for digraph k-vertex connectivity

Algorithm Las Vegas k-connectivity
Input: a digraph $G = (V, E)$.
Output: accept G if it is k-connected, otherwise reject it.
begin
  for $i = 1$ to k **do**
  begin
    Choose a vertex $y_i \in V - \{y_1, \ldots, y_{i-1}\}$;
    loop
      Call *testroot*(y_i, G) and *testroot*$(y_i, \text{rev}(G))$ and
      if both calls of *testroot* succeed, then **exit** the loop;
      choose one of G or rev(G) on which *testroot* fails, and
      let v be a vertex causing *testroot* to fail;
      compute the affine hull H of the embedding of $\Gamma_k(v)$;
      let S be the set of vertices s embedded in H
      such that there is an edge (s, w) with either
      $w = y_i$ or w *not* embedded in H;
      if $|S| < k$ **then** reject G and **stop** ;
    forever;
  end;
  accept G;
end.

Figure 3: Las Vegas algorithm for digraph k-vertex connectivity

connectivity of digraphs can be obtained from the above algorithm by eliminating the possibility of both accepting a digraph that is not k-vertex connected, and rejecting a digraph that is k-vertex connected. To handle the first possibility, we simply choose k distinct vertices $y_1, \ldots, y_k$, and for both the digraphs G and rev(G) call *testroot* k times with the ith call using y_i as the root vertex.

Consider the second possibility. If a call of *testroot* with root vertex, say, y fails for the vertex v, i.e., if we find that the modular X-embedding f computed by this call of *testroot* has rank$(f(\Gamma_k(v))) < k$, then we attempt to find a separator S of cardinality $< k$ whose deletion destroys all paths from v to y. Although there may be many such separators, among the minimum cardinality v-y separators there is a unique one that is "closest" to v, and the affine hull of $f(\Gamma_k(v))$ is determined by this separator. To see this, define a partial order on the minimum cardinality v-y separators such that two separators S_1 and S_2 are related if S_1 intersects every path from v to S_2. It is known that this partial order is a lattice, see [Lo 79]; the unique minimum element of this lattice, denoted $S(v, y)$, is the separator which determines the affine hull of $f(\Gamma_k(v))$. This gives us the next lemma.

LEMMA 5.3. *Let $G = (V, E)$ be a digraph, and v and y be vertices of G. Suppose that $p(v, y) < k$. Let $S = S(v, y)$, and let T be the set of vertices having paths to y in $G - S$. Then, for a random modular X-embedding f, where X is a set of k predecessors of y, with probability at least $1 - 1/n^2$ the affine hull of $f(\Gamma_k(v))$ does not contain any point of $f(T)$.*

When a call of *testroot* with root vertex y fails for the vertex v, we first compute the affine hull H of $\Gamma_k(v)$, and then find the set of vertices S that are embedded in H and either have a successor embedded outside H or have y as a successor. Note that S is a v-y separator. If S has cardinality $< k$, then we have showed that G is not k-vertex connected, otherwise this call of *testroot* has been futile and we repeat the call with the same root vertex. See Figure 3. The above discussion gives us the next theorem.

THEOREM 5.2. *There is a Las Vegas algorithm with an expected running time of $O((M(n) + nM(k)) \cdot k)$ to test the k-vertex connectivity of any digraph. The algorithm accepts a digraph iff it is k-vertex connected.*

Both the algorithms in this section can be implemented on the PRAM model of computation. On an EREW PRAM, both algorithms have running times of $O((\log n)^2)$. The critical computation for both parallel algorithms is to invert the matrix A in equation $(\star)$. Using the algorithms of Kaltofen and Pan [KP 91], with high probability, an $n \times n$ matrix can be inverted on an EREW PRAM in time $O((\log n)^2)$ using $(M(n) \log n)$

processors; for completeness, we cite the main theorem of [KP 91].

THEOREM 5.3. (KALTOFEN AND PAN) *Let n be a number ≥ 1, $\mathcal{K}$ be a field of characteristic zero or $> n$, and $A \in \mathcal{K}^{n\times n}$ be a nonsingular matrix. There exists a randomized algebraic circuit for computing A^{-1} that has size $O(M(n)\log n)$ and depth $O((\log n)^2)$, and that uses $O(n)$ random input elements drawn uniformly from a subset S of $\mathcal{K}$ (besides the n^2 input elements of A). The circuit outputs the n^2 entries of A^{-1} with probability $\geq 1-(3n^2/|S|)$. With probability $\leq 3n^2/|S|$ (or if the matrix A is singular) the circuit divides by zero.*

The next theorem gives the complexity of both our parallel algorithms.

THEOREM 5.4. *The k-vertex connectivity of a digraph can be tested on the EREW PRAM model of computation by :*

1. *A Monte Carlo algorithm in a running time of $O((\log n)^2)$ using $O((M(n)+nM(k))\cdot(\log n)^2)$ processors. If the digraph is k-connected (is not k-connected), then the algorithm accepts the digraph (rejects the digraph) with probability $> 1-(1/n+1/n^2)$.*

2. *A Las Vegas algorithm in an expected running time of $O((\log n)^2)$ using $O((M(n)+nM(k))\cdot k\log n)$ processors. The algorithm accepts a digraph iff it is k-vertex connected.*

References

[B 82] M. Becker, W. Degenhardt, J. Doenhardt, S. Hertel, G. Kaninke, W. Keber, K. Mehlhorn, S. Näher, H. Rohnert, and T. Winter, "A probabilistic algorithm for vertex connectivity of graphs", *Information Processing Letters* **15** (1982), 135–136.

[B 78] B. Bollobàs, *Extremal Graph Theory,* Academic Press, London, 1978.

[CM 88] J. Cheriyan and S. N. Maheshwari, "Finding nonseparating induced cycles and independent spanning trees in 3-connected graphs," *Journal of Algorithms* **9** (1988), 507–537.

[CKT 91] J. Cheriyan, M. Kao and R. Thurimella, "Scan-first search and sparse certificates : An improved parallel algorithm for k-vertex connectivity", preprint, 1991; to appear in the *SIAM J. Computing.*

[CW 87] D. Coppersmith and S. Winograd, "Matrix multiplication via arithmetic progressions", *Proc. 19th Annual ACM Symposium on Theory of Computing,* 1987, 1–6.

[E 75] S. Even, "An Algorithm for Determining Whether the Connectivity of a Graph is at Least k," *SIAM J. Computing* **4** (1975), 393–396.

[E 79] S. Even, *Graph Algorithms,* Computer Science Press, Potomac, Md., 1979.

[G 80] Z. Galil, "Finding the vertex connectivity of graphs," *SIAM J. Computing* **9** (1980), 197–199.

[IR 88] A. Itai and M. Rodeh, "The multi-tree approach to reliability in distributed networks," *Information and Computation* **79** (1988), 43–59.

[KP 91] E. Kaltofen and V. Pan, "Processor efficient parallel solution of linear systems over an abstract field," *Proc. 3rd Annual ACM Symposium on Parallel Algorithms and Architectures,* 1991.

[LEC 66] A. Lempel, S. Even and I. Cederbaum, "An algorithm for planarity testing of graphs." In *Theory of Graphs : Internat. Sympos.: Rome* P. Rosenstiehl, Ed., 215–232, Gordon and Breach, New York, 1966.

[LLW 88] N. Linial, L. Lovász and A. Wigderson, "Rubber bands, convex embeddings and graph connectivity," *Combinatorica* **8** (1988), 91–102.

[Lo 73] L. Lovász, "Connectivity in Digraphs", *J. Combinatorial Theory (B)* **15** (1973), 174–177.

[Lo 79] L. Lovász, *Combinatorial Problems and Exercises,* North-Holland, 1979.

[LSS 89] L. Lovász, M. Saks, and A. Schrijver, "Orthogonal representations and connectivity of graphs", *Linear Algebra and its Applications* **114/115** (1989), 439–454.

[M 84] K. Mehlhorn, *Data Structures and Algorithms 2 : Graph Algorithms and NP-Completeness,* Springer-Verlag, Berlin, 1984.

[MVV 87] K. Mulmuley, U. V. Vazirani, and V. V. Vazirani, "Matching is as easy as matrix inversion", *Combinatorica* **7** (1987), 105–113.

[NI 90] H. Nagamochi and T. Ibaraki, "Linear Time Algorithms for Finding a Sparse k-Connected Spanning Subgraph of a k-Connected Graph", *Algorithmica,* to appear.
"Linear time algorithms for finding k-edge-connected and k-node-connected spanning subgraphs", Technical Report #89006, Dept. of Applied Mathematics & Physics, Faculty of Engineering, Kyoto University, 1989.

[P 87] V. Pan, "Complexity of parallel matrix computations", *Theoret. Comput. Sci.* **54** (1987), 65–85.

[P 91] J. Plehn, Ph. D. Thesis, University of Bonn, Germany.

[S 79] A. Schrijver, "Fractional packing and covering." In *Packing and Covering in Combinatorics,* A. Schrijver, Ed., 201–274, Mathematisch Centrum, Amsterdam, 1979.

[Sc 80] J. T. Schwartz, "Fast probabilistic algorithms for verification of polynomial identities", *J. ACM* **27** (1980), 701–717.

[W 87] R. W. Whitty, "Vertex-disjoint paths and edge-disjoint branchings in directed graphs" *J. Graph Theory* **11** (1987), 349–358.

[ZI 89] A. Zehavi and A. Itai, "Three tree-paths", *J. Graph Theory* **13** (1989), 175–188.

[Z 79] R. E. Zippel, "Probabilistic algorithms for sparse polynomials." In *Proc. EUROSAM 79,* Ng, Ed., Lecture Notes in Computer Science **72**, 216–226, Springer-Verlag, 1979.