

DIRECTED s - t NUMBERINGS, RUBBER BANDS, AND TESTING DIGRAPH k -VERTEX CONNECTIVITY

JOSEPH CHERIYAN and JOHN H. REIF

Received October 3, 1991

Revised November 2, 1992

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 ($M(n) = O(n^{2.376})$). 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 [2], and has diverse applications in computer science, electrical engineering and operations research [7], [18], [11]. 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 $O(\max\{k, \sqrt{n}\}k\sqrt{nm})$ [8]; for $k > \sqrt{n}$, the probabilistic technique of [1] 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 $O(\max\{k, \sqrt{n}\}k^2n^{1.5})$ [20] and

[4]. 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 $O(n^{2.376})$ [5]. Recently, [15] 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 [15] 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 [17].

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 [15]. 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 in general position 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 scheme and using the method of [15], 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 , $k = 2, 3, \dots$, 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 our 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 $k = 1$, the best deterministic parallel algorithm

known runs in $O((\log n)^2)$ time and uses $M(n)$ 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 $O(nmM(n))$ [19]. Our Monte Carlo algorithm improves on the number of processors used by the best previous algorithm by a factor of at least $(nm/(\log n)^2)$ 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 $(n^2m/(\log n)^2)$. Even comparing with the best previous parallel $O((\log n)^3)$ -time algorithm for testing k -vertex connectivity (which is significantly slower), the number of processors used by our Monte Carlo algorithm is much less. The $O((\log n)^3)$ -time algorithm executes at least n copies of the parallel Monte Carlo maximum matching algorithm of [9], see also [13], and uses at least $(n^2 \cdot M(n))$ processors overall.

For undirected graphs, a nondegenerate 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 [14]. Since then this notion has been used for planarity testing [14], fault-tolerant protocols for distributed computers [11], 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 every k -vertex connected directed graph has a nondegenerate 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 [24]; very recently, we have learned of another construction due to Plehn [21].

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. 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 $X \subset V$ 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} .

For any $v \in V$, call the set $\{w : (v, w) \in E\}$ the *successors* of v , and denote this set by $\Gamma(v)$. For a subset X of V , $\Gamma(X)$ denotes $(\cup_{x \in X} \Gamma(x)) - X$, and $\gamma(X)$ denotes $|\Gamma(X)|$. Call the set $\{u : (u, v) \in E\}$ the *predecessors* of v . Define the *reversal* of G , $\text{rev}(G)$, to be the digraph $(V, \{(w, v) : (v, w) \in E\})$.

A *separator* is a subset S of V such that $G - S$ is not strongly connected. For subsets X and Y of V , an *X - Y separator* is a subset S of V such that $G - S$ has no path from $X - S$ to $Y - S$, and $p(X, Y)$ denotes the maximum number of vertex disjoint paths from X to Y (here, the paths must have distinct end-vertices). A set of directed paths with the same start vertex, say x , and the same end vertex, say y ($x \neq y$), is called *internally vertex disjoint* if these paths have no vertices in common except x and y . For two distinct vertices x and y , we use $p(x, y)$ to denote the maximum number of internally vertex disjoint paths from x to y (note the difference from the previous use of p).

A *branching* B is a spanning acyclic subgraph of G such that one vertex, called the *root*, has outdegree zero and all other vertices have outdegree one. 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 ($v \neq z$), the two paths from v to z in the two branchings are internally vertex disjoint. For $k > 2$, k branchings are called *independent* if they are pairwise independent.

For a set of points X in \mathbf{R}^k , $\text{conv}(X)$ denotes the convex hull of X , and $\text{rank}(X)$ 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).

3. s - t numberings for directed graphs

We generalize the notion of s - t numberings to 2-vertex connected digraphs. The numbering π described in the following theorem is called a *directed s - t numbering*.

Theorem 3.1. *A directed graph $G = (V, E)$ with at least three vertices is 2-vertex connected if and only if for every pair of vertices s, t there is a numbering $\pi : V \rightarrow \{1, \dots, n\}$ 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 $\pi(u) < \pi(v)$.*

Proof. To prove the “if part”, first focus on an arbitrary but fixed pair of vertices s and t . Notice that if there is a numbering π satisfying the conditions of the theorem,

then for each vertex $v \in V - \{s, t\}$, G has a path from v to s and a path from v to t such that these paths have no common vertex except v . Since this holds for each pair s, t , G must be 2-vertex connected.

To prove the “only-if part”, we use induction on n to prove the following statement for an arbitrary but fixed pair of vertices s and t : If for each $v \in V - \{s, t\}$, G has a path from v to s and a path from v to t such that these paths have no common vertex except v , then G has a directed s - t numbering π .

The basis of the induction is easily seen to hold, so we proceed with the induction step.

Using a theorem due to Lovász [16], it can be shown that G has a subgraph $G' = (V, E')$ such that the outdegree of both s and t (in G') is zero, and for each $v \in V - \{s, t\}$, the outdegree of v (in G') is two, and G' has a path from v to s and a path from v to t such that these paths have no common vertex except v . Clearly, any valid numbering π of G' is also a valid numbering of G . Focus on the indegree of the vertices in G' . Since G' has at most $2n - 6$ edges with head vertices in $V - \{s, t\}$ (there must be at least two edges entering $\{s, t\}$), G' has a vertex with indegree at most one.

If G' has a vertex v with indegree zero, then it can be seen that $G' - v$ satisfies the induction hypothesis, 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 $G' \setminus (u, v)$. By applying Menger’s theorem, it can be seen that $G' \setminus (u, v)$ satisfies the induction hypothesis, therefore, we can find a numbering π' for $G' \setminus (u, v)$ 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” 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 $v \in V - \{z\}$, the two paths from v to z , in B_1 and B_2 respectively, are internally vertex disjoint. For $k=2$, this gives a simple solution of 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 [22, pp. 235], and later Itai and Rodeh [11] raised a variant of the conjecture for undirected graphs. For $k \leq 3$, the conjecture for undirected graphs has been resolved independently by [3] and [25]. The conjecture remains open for directed graphs for k greater than two and for undirected graphs for k greater than three.

A construction for two independent branchings has been presented before in [24]. Our construction, which is much simpler than the one in [24], 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 3.2 (Whitty). *A digraph is 2-vertex connected iff for every vertex z the digraph has two independent branchings rooted at z .*

4. Convex directed embeddings

Our main result here is that a digraph G is k -vertex connected iff for every subset X of V with $|X| = k$ there exists at least one nondegenerate convex directed X -embedding of G , 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 $c: E \rightarrow \mathbf{R}$ 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, \dots, v_{n-k} and v_{n-k+1}, \dots, v_n be arbitrary linear orderings of $V - X$ and X , respectively. Using c , we can find an embedding $f^{(c)}: V \rightarrow \mathbf{R}^{k-1}$ with the vertices of X at the corners of the unit simplex, i.e., with $f^{(c)}(v_n) = (0, \dots, 0)$ and for $j = 1, \dots, k-1$, $f^{(c)}(v_{n-k+j}) = e_j$, 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 $f_1^{(c)}(v), f_2^{(c)}(v), \dots, f_{k-1}^{(c)}(v)$. (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 $d = 1, \dots, k-1$,

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

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

$$(\star) \quad AF_d = B_d \quad d = 1, \dots, k-1,$$

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

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

B_d and F_d are column vectors of length $n-k$ with the i th entry ($1 \leq i \leq n-k$) of F_d , F_{id} , standing for the d th co-ordinate of v_i , $f_d(v_i)$, and with the i th entry of B_d being given by

$$B_{id} = \begin{cases} c(v_i, v_{n-k+d}) & \text{if } v_i \text{ is a predecessor of } v_{n-k+d} \in X, \\ 0 & \text{otherwise.} \end{cases}$$

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 . The following result, Theorem 4.2, shows that for every convex directed X -embedding f and for every subset U of V , the affine dimension of $f(U) = \{f(w) \mid w \in U\}$ plus one is less than or equal to $p(U, X)$.

Our main result, Theorem 4.3, states something more: There exists an embedding f such that for 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 , $p(U, X)$.

Lemma 4.1. *If G is strongly connected, then the matrix A is nonsingular.*

Proof. The matrix A has the following key property:

For each i , $i = 1, \dots, n-k$,

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

Further, there is at least one row for which the inequality is strict, because at least one vertex in $V - X$, say v_h , has a successor in X , and therefore

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

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

$$Ax = 0,$$

otherwise, A would have full rank.

By considering the inner product of the h th 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

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

In other words, $|x_h| < |x_j|$, for some $j \in \{1, \dots, n-k\}$.

After renumbering the vertices if necessary, it follows that there is an l ($l < n-k$) such that:

$$(4) \quad |x_1| = \dots = |x_l| > |x_{l+1}| \geq \dots \geq |x_{n-k}|.$$

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

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

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

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

This completes the proof. ■

Remark. If G is k -vertex connected for k at least one, then G is strongly connected and by the lemma A is nonsingular. Notice that the hypothesis of the lemma can be weakened without affecting the proof: to ensure that A is nonsingular, the digraph $G \setminus X$ obtained by contracting the vertex set X into a single vertex, say x^* , should have a branching with root x^* .

Theorem 4.2. Let $G = (V, E)$ be a digraph, let X be a subset of V and let k denote $|X|$. Then for every convex directed X -embedding $f: V \rightarrow \mathbf{R}^{k-1}$ of G and for every subset U of V , $U \neq \emptyset$, $\text{rank}(f(U)) \leq p(U, X)$.

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 . Clearly, U is contained in $W \cup S$. For every convex directed X -embedding f , the following argument shows that each extreme point of $\text{conv}(f(W \cup S))$ belongs to $f(S)$. Consider any vertex u in W , and note that by the definition of W , $u \notin X$ and $u \notin S$. Since f is a convex directed embedding it follows that $f(u)$ is contained in $\text{conv}(f(\Gamma(u)))$, consequently, since $\Gamma(u)$ is a subset of $W \cup S - \{u\}$, $f(u)$ is contained in $\text{conv}(f(W \cup S - \{u\}))$. This shows that $\text{conv}(f(W \cup S)) = \text{conv}(f(S))$. It now follows that $\text{rank}(f(U)) \leq \text{rank}(f(W \cup S)) = \text{rank}(f(S)) \leq |S| = p(U, X)$. ■

Theorem 4.3. *Let $G=(V, E)$ be a digraph, let X be a subset of V and let k denote $|X|$. Then G has a convex directed X -embedding $f: V \rightarrow \mathbf{R}^{k-1}$ such that for every subset U of V , $U \neq \emptyset$, $\text{rank}(f(U)) = p(U, X)$.*

Proof. Consider a fixed but arbitrary subset U' of V , and focus on $p(U', X)$. If $p(U', X)$ equals k , then let $U = \{u_1, u_2, \dots, u_k\}$ denote any subset of U' such that G has k vertex disjoint paths from U to X , i.e., the i th path starts with u_i and ends with a distinct vertex of X . Otherwise, if $p(U', X) = k'$ is less than k , then let $U = \{u_1, \dots, u_{k'}\} \cup (X - \{x_1, \dots, x_{k'}\})$ denote any subset of $U' \cup X$ obtained by taking the set of start vertices $\{u_1, \dots, u_{k'}\} \subset U'$ of k' vertex disjoint paths from U' to X , and adding all the vertices of X except the vertices $x_1, \dots, x_{k'}$ that belong to these k' vertex disjoint paths. Let the vertices of U be denoted by v_1, \dots, v_k . Observe that $p(U, X)$ equals k , by the construction of U .

Let $c: E \rightarrow \mathbf{R}$ be a vector of positive coefficients, and $f^{(c)}$ be the embedding obtained by solving equation (*). The rank of $f^{(c)}(U)$ is determined by the matrix $g(c, U)$ given by

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

The rank of $f^{(c)}(U)$ is less than 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, \dots, c_m of c , since each $f_d^{(c)}(u_i)$ is an entry of the matrix $F_d = A^{-1}B_d$ (see equation (*)). 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 there exists a vector of coefficients c' such that for the resulting embedding $f^{(c')}$ the affine rank of $f^{(c')}(U)$ equals $p(U, X)$. Let P_1, \dots, 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, \dots, 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 (in \mathbf{R}^{k-1}) 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, \dots, 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 $(y_1 = u_i), y_2, \dots, y_{l-1}, (y_l = x_i)$. 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 one. As γ increases the equilibrium distance of y_{l-1} and $y_l = x_i$ decreases, and for a sufficiently high γ the distance becomes less than ϵ/n . Repeating the above argument for the other vertices $y_{l-2}, \dots, y_2, (y_1 = u_i)$ 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 equal to $p(U, X) = k$, then $f^{(c')}(U')$ has rank equal to $p(U', X)$.

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 such that for all subsets U' of V $\text{rank}(f^{(c)}(U'))$ equals $p(U', X)$. ■

Theorem 4.4. *Let $G = (V, E)$ be a digraph, and k be an integer, $k = 2, \dots, n - 1$. 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 $f : V \rightarrow \mathbf{R}^{k-1}$ such that for every k -vertex subset Y of V $\text{rank}(f(Y))$ equals $p(Y, X) = k$. 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 at least $\text{rank}(f(Y))$ by Theorem 4.2, and $\text{rank}(f(Y)) = k$ since f is in general position. Consequently, G is k -vertex connected. ■

To gain computational efficiency, instead of finding a directed X -embedding f in \mathbf{R}^{k-1} , we use the [15] method of computing a directed X -embedding f in the $k - 1$ dimensional linear space over a finite field F ; the resulting f is called a *modular directed X -embedding*. The computation is done over the field of integers modulo a prime p , Z_p . A *random modular directed X -embedding* is constructed as follows: Fix a random prime p in the interval $[n^5, n^6]$ and do the computations over $(Z_p)^{k-1}$. Choose a random nonzero coefficient function $c : E \rightarrow (Z_p - \{0\})$ on the edges, and compute a directed X -embedding $f^{(c)} : V \rightarrow (Z_p)^{k-1}$ by solving equation (\star) . Two remarks on computing random modular directed 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 reals. Since p is drawn randomly from the primes in the interval $[n^5, n^6]$, and since $\det(A)$ is nonzero over

the reals, 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 coefficient vector $c \in (Z_p - \{0\})^{|E|}$ gives a modular directed X -embedding $f^{(c)}$ with $\text{rank}(f^{(c)}(U))$ less than $p(U, X)$. Consequently, for a k -vertex connected digraph G , it is *not* necessarily true that there exists a modular directed X -embedding f such that *every* subset U of V with cardinality k has $\text{rank}(f(U)) = k$.

Lemma 4.5. *Let $G = (V, E)$ be a digraph, let X be a subset of V and let k denote $|X|$. Then for any fixed subset U of $V - X$, a random modular directed X -embedding $f: V \rightarrow (Z_p)^{k-1}$ satisfies $\text{rank}(f(U)) = p(U, X)$ with probability at least $1 - (1/n^3)$.*

Proof. Consider a vector c drawn at random from $(Z_p - \{0\})^{|E|}$. With probability greater than $1 - \frac{1}{n^4}$, the matrix A in equation (\star) has $\det(A) \neq 0 \pmod p$. Moreover, assuming that $\det(A) \neq 0 \pmod p$, the determinant $\det(g(c, U))$ in the proof of Theorem 4.3 is a rational function of degree at most $2k(n - k) \leq \frac{1}{2}n^2$ whose denominator is nonzero, because each entry of the matrix $g(c, U)$ is a rational function of degree at most $1 + (n - k) \leq 2(n - k)$ whose denominator (namely, $\det(A)$) is nonzero. Note that each entry is obtained by solving equation (\star) . Now applying the Zippel-Schwartz lemma [26], [23], the probability that $\det(g(c, U)) = 0 \pmod p$ (assuming that $\det(A) \neq 0 \pmod p$) is at most $\frac{n^2}{2p}$. Therefore, the probability that $\text{rank}(f^{(c)}(U))$ is less than k , which is at most the probability that either $\det(A) = 0 \pmod p$ or $\det(g(c, U)) = 0 \pmod p$, is at most $\frac{1}{n^4} + \frac{n^2}{2p} \leq \frac{1}{n^3}$. ■

5. Algorithms for digraph k -vertex connectivity

Subroutine *test root*(z, G)

Input: Digraph $G = (V, E)$, and a “root vertex” z .

Output: If the algorithm returns success, then G has k internally vertex disjoint paths from each $v \in V - \{z\}$ to z . If the algorithm returns failure, then with probability $> 1 - 1/n^2$ G has a separator of size $< k$.

begin

 choose a random prime $p \in [n^5, n^6]$ and a random vector $c \in (Z_p - \{0\})^{|E|}$;

 let X be a set of k predecessors of z ;

 compute the modular directed X -embedding $f^{(c)}$ by solving equation (\star) ;

for all $v \in V - \{z\}$ **do**

 compute the rank of $f^{(c)}(\Gamma_k(v)) = \{f^{(c)}(w) \mid w \in \Gamma_k(v)\}$;

if matrix A is singular over Z_p or some v computes a rank $< k$

then return failure

else return success;

end.

Fig. 1. Constructing and testing a modular directed X -embedding

* The number of distinct prime divisors of $\det(A)$ is $\leq \log |\det(A)| \leq \log(np)^n \leq 7n \log n$; and there are $\Omega(n^6/(\ln n))$ primes in the interval $[n^5, n^6]$, by the prime number theorem [10].

Algorithm Monte Carlo k -connectivity

Input: Digraph $G=(V,E)$.

Output: If G is k -connected, accept it with probability $>1-1/n$, otherwise, reject it with probability $>1-1/n^2$.

begin

 let $k'=\lceil \log n/\log(n/k) \rceil$;

for $i=1$ to k' **do**

begin

 choose a *random* vertex $y_i \in V$;

 call *test root*(y_i, G) and *test root*($y_i, \text{rev}(G)$) and

 if either call of *test root* fails, then reject G and **stop**;

end;

 accept G ;

end.

Fig. 2. Monte Carlo algorithm for digraph k -vertex connectivity

Our algorithms are similar to those of [15]. 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, *test root*, has two inputs: a digraph $G=(V,E)$ and a “root vertex” $z \in V$. The subroutine tests whether G has k internally vertex disjoint paths from each $v \in V-\{z\}$ to z , or whether G has a separator of size less than k . Both outcomes are possible for the same G , since every separator of size less than k may contain z .

Let X be a fixed but arbitrary subset of the predecessors of z with $|X|=k$. For every vertex $v \in V-X$, 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 $|\Gamma_k(v)|=k$. For vertices $v \in X$, let $\Gamma_k(v)$ denote the union of $\{v\}$ and a fixed but arbitrarily chosen set of $k-1$ vertices from $\Gamma(v)-\{z\}$. Consider any vertex $v \in V-\{z\}$. To efficiently test for k internally vertex disjoint paths from v to z , it suffices to check whether G has k vertex disjoint paths from $\Gamma_k(v)$ to X . If G has these paths, then it also has k internally vertex disjoint paths from v to z . On the other hand, if G does not have k vertex disjoint paths from $\Gamma_k(v)$ to X , then by Menger’s theorem G must have a separator of size less than k .

The subroutine *test root*, given a root vertex z , works as follows: First, a random modular directed X -embedding $f:V \rightarrow (Z_p)^{k-1}$ is computed by choosing a random prime p and random nonzero coefficients for the edges of G (i.e., choosing a random vector $c \in (Z_p-\{0\})^{|E|}$), and solving the linear system (\star) . Then for every $v \in V-\{z\}$, we test whether $\text{rank}(f(\Gamma_k(v)))$ is equal to k . If the matrix A is singular over Z_p or if some vertex v fails the test, then the call of *test root* fails, otherwise, the call succeeds. See Figure 1.

Coming to the running time analysis, the linear system (\star) has to be solved in order to construct a modular directed embedding: We first compute the inverse of A in time $O(M(n-k))$, and then for $d=1, \dots, k$ “simultaneously” compute $F_d=f_d:(V-X) \rightarrow Z_p$, by multiplying A^{-1} and the matrix $(B_1 \dots B_k)$ 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 $O(M(\min\{k, n-k\}))$.

Lemma 5.1. *If test root returns success, then for each $v \in V - \{z\}$, the digraph has k internally vertex disjoint paths from v to z , and if test root returns failure, then with probability $> 1 - 1/n^2$ the digraph has a separator of size less than k . The running time of test root is $O(M(n-k) \cdot \lceil k/(n-k) \rceil + n \cdot M(\min\{k, n-k\}))$.*

The following well known result [6, 7] shows how test root can be used to test for k -vertex connectivity.

Lemma 5.2. *Let z_1, \dots, 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 $\text{rev}(G)$, test root succeeds with each of the k root vertices $z = z_1, \dots, z_k$.*

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 less than $1/n$, then we can improve on the number of calls to test root by fixing an appropriate value k' less than k and calling test root for k' randomly chosen root vertices $y_1, \dots, y_{k'}$. An execution of the algorithm may be erroneous: If G is not k -vertex connected, but every separator of G with size less than k contains the randomly chosen root vertices $y_1, \dots, y_{k'}$, then the algorithm would accept G . To fix the value of k' , note that the probability of an erroneous result is at most $((k-1)/n)^{k'}$. This idea has been used before by [1], [18] and [15].

Lemma 5.3. *Let k' be an integer such that $((k-1)/n)^{k'} < 1/n$; k' may be taken to be $\lceil \log n / \log(n/k) \rceil$. Let $y_1, \dots, y_{k'}$ be k' randomly chosen vertices of the digraph G . If for both the digraphs G and $\text{rev}(G)$, test root succeeds with each of the k' root vertices $z = y_1, \dots, y_{k'}$, then with probability $> 1 - (1/n)$ G is k -vertex connected.*

As mentioned before, with probability less than $1/n$ the Monte Carlo algorithm may accept a digraph that is not k -vertex connected. Also, with probability less than $1/n^2$ the algorithm may reject a digraph that is k -vertex connected, because according to Lemma 4.5 with probability less than $1/n^2$ either the matrix A may be singular over Z_p or the embedding f computed by the algorithm may have a vertex v with $\text{rank}(f(\Gamma_k(v))) < k$. See Figure 2.

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 5.4. *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 -vertex connected (is not k -vertex connected), then the algorithm accepts the digraph (rejects the digraph) with probability $> 1 - 1/n$.*

A Las Vegas algorithm for testing the k -vertex 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 choose k distinct vertices y_1, \dots, y_k , and for both the digraphs G and $\text{rev}(G)$ call test root k times with the i th call using y_i as the root vertex.

Consider the second possibility. If a call of test root with root vertex y and k -subset of y 's predecessors X fails for the vertex v , i.e., if we find that the modular directed X -embedding f computed by this call of test root has $\text{rank}(f(\Gamma_k(v))) <$

Algorithm Las Vegas k -connectivity

Input: Digraph $G=(V, E)$.

Output: Accept G if it is k -vertex connected, otherwise reject it.

begin

for $i=1$ to k **do**

begin

 choose a vertex $y_i \in V - \{y_1, \dots, y_{i-1}\}$;

loop

 call $test_root(y_i, G)$ and $test_root(y_i, rev(G))$ and

 if both calls of $test_root$ succeed, then **exit** the loop;

 (if $test_root$ fails, assume it returns the embedding $f^{(c)}$
 and a vertex v causing failure;)

 choose one of G or $rev(G)$ on which $test_root$ fails, and

 let v be a vertex causing $test_root$ 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 either $s \in X$ or s has a successor w

 which is *not* embedded in H ;

 if $|S| < k$ then reject G and **stop** ;

forever;

end;

 accept G ;

end.

Fig. 3. Las Vegas algorithm for digraph k -vertex connectivity

k , then we attempt to find a separator S of cardinality less than k whose deletion destroys all paths from $\Gamma_k(v)$ to X . Suppose that $p(\Gamma_k(v), X) = k'$ is less than k . For any $\Gamma_k(v)$ - X separator S with cardinality k' let $Q(S)$ denote the maximum subset of $V - S$ such that $G - S$ has no path from $Q(S)$ to $X - S$. The following result is well known, and may be proved using the submodularity of γ .

Lemma 5.5. *Let S^* be a $\Gamma_k(v)$ - X separator with cardinality $p(\Gamma_k(v), X) = k'$ such that $Q(S^*)$ is maximal over all such separators S . Then S^* is unique, and for each vertex $v \in V - (S^* \cup Q(S^*))$ the number of vertex disjoint paths from $\{v\} \cup S^*$ to X , $p(\{v\} \cup S^*, X)$, is greater than k' .*

Combining this result with Lemma 4.5, and noting that the affine hull of $f(S^*)$ contains the affine hull of $f(\Gamma_k(v))$ gives the key lemma for analyzing our Las Vegas algorithm.

Lemma 5.6. *Suppose that there is a vertex $v \in V - \{z\}$ with $p(\Gamma_k(v), X)$ less than k , and let S^* and $Q(S^*)$ be as above. For a random modular directed X -embedding f with probability at least $1 - 1/n^2$ no vertex of $V - (S^* \cup Q(S^*))$ is embedded in the affine hull of $f(\Gamma_k(v))$.*

When a call of $test_root$ 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 that either belong to X or have a successor embedded outside H . Note that S is a $\Gamma_k(v)$ - X separator. If S has cardinality less than k , then we have showed

that G is not k -vertex connected, otherwise this call of *test_root* 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 5.7. *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 [12], 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 [12].

Theorem 5.8 (Kaltofen, 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 \mathcal{S} of \mathcal{K} (besides the n^2 input elements of A). The circuit outputs the n^2 entries of A^{-1} with probability $\geq 1 - (3n^2/|\mathcal{S}|)$. With probability $\leq 3n^2/|\mathcal{S}|$ (or if the matrix A is singular) the circuit divides by zero.*

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

Theorem 5.9. *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 $(M(n) + nM(k)) \cdot (\log n)^2$ processors. If the digraph is k -vertex connected (is not k -vertex 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 $(M(n) + nM(k)) \cdot k \log n$ processors. The algorithm accepts a digraph iff it is k -vertex connected.*

Acknowledgments. We thank the referees for their comments.

References

- [1] 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.
- [2] B. BOLLOBÁS: *Extremal Graph Theory*, Academic Press, London, 1978.
- [3] 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.

- [4] J. CHERIYAN, M. KAO, and R. THURIMELLA: Scan-first search and sparse certificates: An improved parallel algorithm for k -vertex connectivity, *SIAM J. Computing* **22** (1993), 157–174.
- [5] D. COPPERSMITH, and S. WINOGRAD: Matrix multiplication via arithmetic progressions, *J. Symbolic Comp.* **9** (1990), 23–52.
- [6] S. EVEN: An algorithm for determining whether the connectivity of a graph is at least k , *SIAM J. Computing* **4** (1975), 393–396.
- [7] S. EVEN: *Graph Algorithms*, Computer Science Press, Potomac, Md., 1979.
- [8] Z. GALIL: Finding the vertex connectivity of graphs, *SIAM J. Computing* **9** (1980), 197–199.
- [9] Z. GALIL, and V. PAN: Improved processor bounds for combinatorial problems in RNC, *Combinatorica* **8** (1988), 189–200.
- [10] G. H. HARDY, and E. M. WRIGHT: *An Introduction to the Theory of Numbers*, Oxford Univ. Press, Oxford, 1979.
- [11] A. ITAI and M. RODEH: The multi-tree approach to reliability in distributed networks, *Information and Computation* **79** (1988), 43–59.
- [12] 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.
- [13] R. M. KARP, E. UPFAL and A. WIGDERSON: Constructing a perfect matching is in random NC, *Combinatorica* **6** (1986), 35–48.
- [14] 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.
- [15] N. LINIAL, L. LOVÁSZ and A. WIGDERSON: Rubber bands, convex embeddings and graph connectivity, *Combinatorica* **8** (1988), 91–102.
- [16] L. LOVÁSZ: Connectivity in digraphs, *J. Combinatorial Theory (B)* **15** (1973), 174–177.
- [17] L. LOVÁSZ, M. SAKS, and A. SCHRIJVER: Orthogonal representations and connectivity of graphs, *Linear Algebra and its Applications* **114/115** (1989), 439–454.
- [18] K. MEHLHORN: *Data Structures and Algorithms 2: Graph Algorithms and NP-Completeness*, Springer-Verlag, Berlin, 1984.
- [19] K. MULMULEY, U. V. VAZIRANI, and V. V. VAZIRANI: Matching is as easy as matrix inversion, *Combinatorica* **7** (1987), 105–113.
- [20] H. NAGAMACHI and T. IBARAKI: 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. *Algorithmica* **7** (1992), 583–596.
- [21] J. PLEHN: Ph. D. Thesis, University of Bonn, Bonn, Germany.
- [22] A. SCHRIJVER: Fractional packing and covering, in: *Packing and Covering in Combinatorics*, A. Schrijver, Ed., 201–274, Mathematisch Centrum, Amsterdam, 1979.
- [23] J. T. SCHWARTZ: Fast probabilistic algorithms for verification of polynomial identities, *J. ACM* **27** (1980), 701–717.
- [24] R. W. WHITTY: Vertex-disjoint paths and edge-disjoint branchings in directed graphs, *J. Graph Theory* **11** (1987), 349–358.
- [25] A. ZEHAVI, and A. ITAI: Three tree-paths, *J. Graph Theory* **13** (1989), 175–188.

- [26] R. E. ZIPPEL: Probabilistic algorithms for sparse polynomials, in: *Proc. EUROSAM 79*, Lecture Notes in Computer Science **72**, 216–226, Springer-Verlag, 1979.

Joseph Cheriyan

*Department of
Combinatorics & Optimization
University of Waterloo
Waterloo, Ontario
Canada N2L 3G1
jcheriyan@watdragon.uwaterloo.ca*

John H. Reif

*Department of Computer Science
Duke University
Durham, NC 27706
U.S.A.
reif@cs.duke.edu*