COMBINATORICA
Akadémiai Kiadó — Springer-Verlag

RUBBER BANDS, CONVEX EMBEDDINGS AND GRAPH CONNECTIVITY

N. LINIAL, L. LOVÁSZ and A. WIGDERSON

Received September 15, 1986
Revised March 17, 1987

We give various characterizations of k-vertex connected graphs by geometric, algebraic, and "physical" properties. As an example, a graph G is k-connected if and only if, specifying any k vertices of G, the vertices of G can be represented by points of $\mathbf{R}^{k-1}$ so that no k are on a hyperplane and each vertex is in the convex hull of its neighbors, except for the k specified vertices. The proof of this theorem appeals to physics. The embedding is found by letting the edges of the graph behave like ideal springs and letting its vertices settle in equilibrium.

As an algorithmic application of our results we give probabilistic (Monte-Carlo and Las Vegas) algorithms for computing the connectivity of a graph. Our algorithms are faster than the best known (deterministic) connectivity algorithms for all $k \geqq \sqrt{n}$, and for very dense graphs the Monte Carlo algorithm is faster by a linear factor.

0. Introduction

The property of k-connectivity of a graph is well-characterized: it is "easy" to exhibit if a graph is not k-connected and, also, if it is. But there is some asymmetry in this: to exhibit that a graph is not k-connected, it suffices to present a separating set with less than k vertices; to exhibit that it is k-connected, we have to present k openly disjoint paths *for each pair of vertices.* Is there a more compact "proof" of k-connectivity, say, an additional structure whose presence gives a trivially checkable proof of k-connectivity? For $k=1$, a spanning tree provides a trivial answer. For $k=2$, various versions of "ear-decompositions" (see, e.g., [10]) give rise to such "proofs". Another structure characterizing 2-connectivity, closely related to ear-decompositions, is an $s-t$ numbering for an edge st: a linear ordering of the nodes, starting with s and ending with t, such that every other node has a neighbor to its left as well as one to its right.

In this paper, we offer some new characterizations of graph k-connectivity, based on geometric and physical intuition. Our main theorem is a geometric characterization of k-vertex connected graphs, generalizing $s-t$ numberings. It says that a graph G is k-connected if and only if G has certain "nondegenerate convex embeddings" in $\mathbf{R}^{k-1}$.

The proof of this theorem appeals to physics. The embedding is found by letting the edges of the graph behave like ideal springs and letting its vertices settle

AMS subject classification (1980): 05 C 40, 52 A 20.

in equilibrium. Algebraic properties of this equilibrium ensure that the embedding it defines is nondegenerate exactly when the graph is k-connected.

We prove a related, purely matrix-theoretical characterization of connectivity. This result is in fact easily derivable from the results of [12] and [9], which give a linear representation of certain matroids arising in the study of graph connectivity, called strict *gammoids*.

These results provide not only more compact characterizations but also new algorithms. We give probabilistic algorithms for computing the connectivity of a graph. The first is a Monte Carlo algorithm that runs in time $O(n^{2.5}+nk^{2.5})$ where n is the number of vertices and k is the vertex connectivity of the input graph. The second is a Las Vegas algorithm (i.e., never errs) that runs in expected time $O(kn^{2.5}+nk^{3.5})$. For comparison, the best known algorithm (which is deterministic!) runs in time $k^3 n^{1.5}+k^2n^2$ [8]. Observe that our algorithms are faster for all $k \geqq \sqrt{n}$, and for very dense graphs the Monte Carlo algorithm is faster by a linear factor!

We also desribe parallel implementations of our algorithms which are substantially more efficient than previous parallel algorithms for graph connectivity.

1. Notation

Graph Theory. Let $G(V, E)$ be a graph. For a vertex $v \in V$, $N(v)=\{u: (v, u) \in E\}$ denotes the neighborhood of v, and $\bar{N}(v)=N(v) \cup \{v\}$. Let X, Y be any two subsets of V. By $p(X, Y)$ we denote the maximum number of vertex disjoint paths from X to Y (disjointness includes the end points!). We say that X and Y are *linked* if $|X|=|Y|=p(X, Y)$. By Menger's Theorem, this is equivalent to saying that no set of fewer than $|X|=|Y|$ vertices covers all $X-Y$ paths in G. The graph G is *k-connected* if $|V|>k$ and any two k-subsets are linked. The largest k for which this holds is the *vertex-connectivity* of G, denoted $k(G)$. It is known that a graph G is k-connected if and only if $|V|>k$ and any two nodes of G are connected by k openly disjoint paths.

Algebra. Let F be any field and $d \geqq 0$. We denote by F^d the d-dimensional linear space over F. Let $X=\{x_1, \ldots, x_m\}$ be a finite set of points in $\mathbf{R}^d$. The *affine hull* $\operatorname{aff}(X)$ of X is the set of all points $\sum_{i=1}^{m} \lambda_i x_i$ with $\sum \lambda_i = 1$. The (affine) *rank* of X is defined by $\operatorname{rank}(X)=1+\dim(\operatorname{aff}(X))$. X is in *general position* if $\operatorname{rank}(Y)=$ $=d+1$ for every $(d+1)$-subset $Y \subseteq X$. If X is not in general position, we call it *degenerate*.

If $F=\mathbf{R}$ then we will also consider the *convex hull* $\operatorname{conv}(X)$ of X. Note that $\operatorname{aff}(\operatorname{conv}(X))=\operatorname{aff}(X)$.

2. Convex embeddings

Our main tool is the following notion of embedding graphs in real linear spaces, which may be interesting for purposes other than the study of connectivity.

Definition 2.1. Let G be a graph and $X \subset V$. A *convex X-embedding* of G is any mapping $f\colon V \to \mathbf{R}^{|X|-1}$ such that for each $v \in V \setminus X$, $f(v) \in \operatorname{conv}(f(N(v)))$. We say that the convex embedding is *in general position* if the set $f(V)$ of points is in general position.

Let us state our main theorem right away.

Theorem 2.2 *Let G be a graph on n vertices and $1<k<n$. Then the following two conditions are equivalent:*

(1) G is k-connected.

(2) For every $X\subset V$ with $|X|=k$, G has a convex X-embedding in general position.

Note that the special case $k=2$ of our theorem asserts the existence of an $s-t$ numbering of a 2-connected graph (see [6]).

Theorem 2.2 will follow from Theorems 2.3 and 2.4 below.

Theorem 2.3. *Let G be a graph and $X\subset V$. Then for every convex X-embedding f of G and every subset $U\subseteq V$ $U\neq\emptyset$,* $\operatorname{rank}(f(U))\leqq p(U, X)$.

Theorem 2.4. *Let G be a graph and $X\subset V$. Then G has a convex X-embedding f such that for every $U\subseteq V$, $U=\emptyset$,* $\operatorname{rank}(f(U))=p(U, X)$.

Proof of Theorem 2.2. (2)→(1). Let X, Y be arbitrary k-subsets of V, and let f be an X-embedding guaranteed by (2). Then by Theorem 2.3, $p(X, Y)\geqq\operatorname{rank}(f(Y))=k$, since $f(V)$ are in general position. Therefore G is k-connected.

(1)→(2). Assume G is k-connected and fix a k-subset $X\subseteq V$. Then Theorem 2.4 implies the existence of a convex X-embedding such that for every k-subset $Y\subseteq V$, $\operatorname{rank}(f(Y))=p(X, Y)\geqq k$, so every k points are in general position. ∎

Proof of Theorem 2.3. Let f be a convex X-embedding and fix a subset $U\subseteq V$. Let $p(U, X)=k$. Then by Menger's Theorem, there is a k-subset $S\subseteq V$ such that $V\setminus S$ contains no (X, U) paths. Let W be the union of connected components of $G\setminus S$ containing a vertex from U. We claim that $f(W)\subseteq\operatorname{conv}(f(S))$. Note that this implies $\operatorname{rank}(f(U))\leqq\operatorname{rank}(f(W\cup S))=\operatorname{rank}(f(S))\leqq|S|=k=p(U, X)$.

To prove the claim, let $u\in W$. Hence $u\notin X$. Since f is a convex X-embedding, $f(u)\in\operatorname{conv}(f(N(u)))\subseteq\operatorname{conv}(f(W\cup S\setminus u))$, so $f(u)$ cannot be an extreme point of $f(W\cup S)$. Hence the only extreme points in $f(W\cup S)$ are members of $f(S)$, i.e., $f(W)\subseteq\operatorname{conv}(f(S))$. ∎

Proof of Theorem 2.4. Let X be given and $|X|=k$. The intuition behind the proof is of a physical nature. Assume that the edges of G are made of ideal rubber bands. Glue the vertices of X to the extremes of a k-simplex in $\mathbf{R}^{k-1}$, and let the remaining vertices settle in a minimum energy equilibrium. It should be clear that if the potential carried by each rubber band is positive then such an equilibrium exists, and furthermore, it is a convex X-embedding. To achieve the non-degeneracy properties required by the theorem, we use a quadratic potential function (namely the rubber bands satisfy Hooke's Law) and exploit our freedom in choosing the elasticity parameters (e.g. the thickness of the rubber bands). For a hystorical survey and in-depth study of the potential function on similar frameworks, see [5].

We proceed formally. Let $X=\{x_0, x_1, \dots, x_{k-1}\}$. Let e_0 be the zero vector and let e_i, $1\leqq i\leqq k-1$ be the i-th unit vector in $\mathbf{R}^{k-1}$. An embedding $g\colon V\to\mathbf{R}^{k-1}$ such that $g(x_i)=e_i$ for $0\leqq i\leqq k-1$ is called an X-embedding. (Such an embedding is not necessarily convex!)

Assign to every edge $(u, v)\in E$ a positive elasticity coefficient c_{uv}, and let $c\in\mathbf{R}^E$ be the vector of coefficients.

Now we can define the potential $P(g, c)$ of any X-embedding (not necessarily convex) g and coefficient vector c. The edge (rubber band) $(u, v)\in E$ carries potential $c_{uv}\|g(u)-g(v)\|^2$ (our norm is Euclidean). Hence the total potential is defined by

$$P(g, c) = \sum_{(u, v)\in E} c_{uv}\|g(u)-g(v)\|^2.$$

Let $f=f_c$ be the embedding for which $P(g, c)$ is minimized. Notice that f is uniquely determined since $P(g, c)$ is a strictly convex function of g. Then f satisfies the equilibrium condition $\frac{dP}{dg}=0$. This is a homogeneous linear system:

$$\sum_{(u, v)\in E} c_{uv}(g(u)-g(v)) = 0 \quad \text{for all} \quad v\in V\setminus X.$$

From the strict convexity of P it follows that this system has a unique solution. Putting $c_v=\sum_{u\in N(v)} c_{uv}$ we see that for all $v\in V\setminus X$,

$$f(v) = \frac{1}{c_v}\sum_{u\in N(v)} c_{uv} f(u),$$

which expresses $f(v)$ as a convex combination of $f(N(v))$. Since f is an X-embedding (i.e., $f(x_i)=e_i$, $0\leqq i\leqq k-1$), we get that f is a convex X-embedding of G.

Now fix a subset $U\subseteq V$, and let $p(U, X)=m$. If $m=0$, we are done by Theorem 2, so assume $m\geqq 1$. For a given vector c, f_c may not satisfy rank $(f_c(U))=m$. However, we will show that this happens only for a set of measure zero of possible vectors c. Since there are only finitely many subsets $U\subseteq V$, the theorem will follow.

Let us consider in detail the set of equations which determine f_c from c. Define $T=(t_{uv})$ $(u, v\in V)$ to be the following symmetric matrix (which was also used by [15]:

$$t_{uv} = \begin{cases} c_{uv} & \text{if} \quad uv\in E, \\ -c_v & \text{if} \quad u = v, \\ 0 & \text{otherwise.} \end{cases}$$

Assume that the vertices in V numbered such that $\{x_0, x_1, \dots, x_{k-1}\}$ appear first, and arrange T as a block matrix of the form

$$T=\begin{bmatrix} T_1 & T_2^T \\ T_2 & T_3 \end{bmatrix}$$

where T_1 is a $k\times k$ matrix and T_3 is an $(n-k)\times(n-k)$ matrix. Let F be an $|V|\times(k-1)$ matrix whose (v, j)-th entry is the j-th coordinate in f_c of the vertex

$v \in V$. Then the condition $f_c(x_i)=e_i$, $0 \leqq i \leqq k-1$ implies

$$F=\begin{bmatrix} 0 & 0 \dots & 0 \\ 1 & 0 \dots & 0 \\ \vdots & \vdots & \vdots \\ 0 & 0 & 1 \\ \hdashline & F_1 & \end{bmatrix}$$

and the stability condition becomes $[T_2 : T_3]F=0$, or equivalently

$$T_3 F_1 = B, \tag{$*$}$$

where B arises from T_2 by dropping its first column. This equation expresses the coordinates of the vertices in $V \setminus X$ as rational functions of the coefficients c. (Recall that $(*)$ has a unique solution since it is the minimum of a strictly convex function.)

Now consider a subset $M \subseteq U$ with $M=\{u_1, u_2, \dots, u_m\}$ and $p(M, X)= =p(U, X)=m$. We want to show that $f_c(M)$ has rank m almost always. The $(m-1)$-dimensional volume of $\operatorname{conv}(f_c(M))$ in $\mathbf{R}^{k-1}$ is given by the determinant $D_c=(\det(AA^T))^{1/2}$, where the matrix A is

$$A=\begin{bmatrix} 1 & f(u_1)_1 & f(u_1)_2 \dots & f(u_1)_{k-1} \\ 1 & f(u_2)_1 & f(u_2)_2 \dots & f(u_2)_{k-1} \\ \vdots & \vdots & \vdots & \vdots \\ 1 & f(u_m)_1 & f(u_m)_2 \dots & f(u_m)_{k-1} \end{bmatrix}.$$

This determinant vanishes exactly when $f_c(M)$ does not have full rank. As the $f_c(u_j)$ are rational functions of the entries of c, the determinant D is either identically zero, or it vanishes only for a set of vectors c of measure zero. We want to exclude the first possibility.

As $p(M, X)=m$, there are m vertex disjoint paths $P_1, P_2, \dots, P_m$ from M to X, and assume without loss of generality that P_i connects u_i to x_i.

Set $E'=\{\cup E(P_i)\colon i=1, \dots, m\}$, and let $c=c(R)$ be defined by $c_e=1$ for all $e \in E \setminus E'$ and $c_e=R$ for all $e \in E'$. We claim that if we let R tend to infinity, then the distances $\|f(u_i)-f(x_i)\|$ will tend to zero. Once they are small enough (all less then $1/(2\sqrt{k})$), the set $f_c(M)$ must have rank m.

Let $f=f_{c(R)}$ and recall that f minimizes the potential $P(g, c(R))$ over all X-embeddings g. Let f' be the embedding with $f'(V(P_i))=e_i$, $1 \leqq i \leqq m$ and (say) $f'(v)=e_0$ for any vertex not in X or $\cup P_i$. Then

$$P(f, c(R)) \leqq P(f', c(R)) \leqq \sqrt{2}\,|E|,$$

as every edge in f' has length at most $\sqrt{2}$, and the coefficients of these edges are 1.

On the other hand,

$$P(f, c(R)) = \sum_{(u,v) \in E} c_{uv} \|f(u)-f(v)\|^2 \geqq \sum_{i=1}^{m} \sum_{(u,v) \in P_i} R \|f(u)-f(v)\|^2.$$

By the Cauchy—Schwartz inequality and the fact that $|P_i|\leqq n$ we get

$$\sum_{i=1}^{m}\sum_{(u,v)\in P_i} R\|f(u)-f(v)\|^2 \geqq \sum_{i=1}^{m}\frac{R}{|P_i|}\Big(\sum_{(u,v)\in P_i}\|f(u)-f(v)\|\Big)^2 \geqq$$
$$\geqq \frac{R}{n}\sum_{i=1}^{m}\|f(x_i)-f(u_i)\|^2 \geqq \frac{R}{n}\max_i\|f(x_i)-f(u_i)\|^2.$$

This proves the claim and thereby the theorem. ▌

3. A related result on matrices

Let us modify the matrix T considered in Section 2 as follows. Let us introduce a variable x_{vv} for each vertex v, and a variable x_{uv} for each edge $uv\in E(G)$. Define $x_{uv}=0$ if u and v are non-adjacent vertices. Let us call the matrix $A_G=(x_{uv})_{u,v\in V(G)}$ the *free adjacency matrix* of G.

A closely related, although not equivalent, result which gives a linear representation of the so-called strict gammoids follows from [12] and [9].

Theorem 3.1. *Let G be a graph, $X, Y\subseteq V(G)$, $|X|=|Y|$, and let A_{XY} be the matrix obtained from A_G by deleting the rows corresponding to X and the columns corresponding to Y. Then* $\det(A_{XY})$ *is not identically* 0 *if and only if X and Y are linked.*

(Note that $\det(A_{XY})$ is a polynomial in the variables x_{uv}.)

Proof. Let $k=|X|=|Y|$. Assume first that X and Y are not linked. Then by Menger's Theorem, G can be written as the union of two graphs G_1 and G_2 such that $|V(G_1)\cap V(G_2)|\leqq k-1$, $X\subseteq V(G_1)$ and $Y\subseteq V(G_2)$. This means that the matrix A_G has the following form:

$$A_G=\begin{bmatrix} A_{11} & A_{12} & 0 \\ A_{21} & A_{22} & A_{23} \\ 0 & A_{32} & A_{33} \end{bmatrix}$$

where there are, say, a, $k-1$, b columns in the first, second and third block of columns, respectively, and similarly for the rows. Furthermore, the rows in X belong to the first two blocks while the columns in Y belong to the last two blocks. So A_{XY} contains as entries all the 0's in the lower left $a\times b$ block. Since $a+b=n-(k-1)>n-k$, the order of A_{XY}, this implies that A_{XY} is singular for any values of the x_{uv}.

Conversely, assume that X and Y are linked. To show that $\det(A_{XY})$ is not identically 0 we exhibit a special choice of the variables for which this determinant is not 0. Let $P_1, \dots, P_k$ be k vertex-disjoint paths linking X and Y. Then A_G has the following form:

$$A_G=\begin{bmatrix} A_1 & 0 & \dots & & 0 \\ 0 & A_2 & \dots & 0 & 0 \\ \vdots & & & \vdots & \vdots \\ 0 & 0 & \dots & A_k & 0 \\ 0 & 0 & \dots & 0 & I \end{bmatrix}$$

where A_i is formed by the rows and columns corresponding to vertices of P_i, and has the form

$$A_i = \begin{bmatrix} 1 & 1 & & & 0 \\ 1 & 1 & 1 & & 0 \\ & 1 & 1 & & \\ & & \ddots & & \\ & & & \ddots & 1 \\ 0 & & & 1 & 1 \end{bmatrix}$$

and I is an identity matrix. Moreover, X consists of the first rows of the first k blocks, and Y consists of the last columns of the first k blocks. Hence A_{XY} is an upper triangular matrix with 1's in the main diagonal, which shows that $\det(A_{XY})=1$. ▮

Remarks. 1. The matrix T occurring in the proof of Theorem 2.4 arises from A_G by the substitution $x_{uv}=c_{uv}$ if uv is an edge and $x_{vv}=-c_v=-\sum_w c_{wv}$. Hence the "only if" part of Theorem 3.1 remains valid if we replace A_G by T. This is not true for the "if" part. The sets $X=Y=\emptyset$ give an obvious counterexample; a less obvious counterexample is given by $X=\{x\}$ and $Y=\{y\}$ where G is disconnected and x and y are distinct nodes of the same connected component of G.

2. The result of Theorem 3.1 is valid over any field. It also remains true if we do not assume that the matrix is symmetric.

4. Algorithmic applications

Our model of computation is the logarithmic cost RAM, see [1]. Both Theorems 2.4 and 3.1 lend themselves to (randomized) graph connectivity algorithms which have good running times and are easily parallelizable.

Using our previous results, we do not directly obtain a test for the k-connectivity of a graph, but rather a test for checking whether or not two given k-tuples are linked, or a test whether or not a given k-tuple is linked to all other k-tuples. The following remarks show how to use such a subroutine in connectivity testing.

1) Let, for every vertex $v\in V$, $N_k(v)$ denote an arbitrary k-subset of $N(v)$. Then G is k-connected iff $N_k(u)$ and $N_k(v)$ are linked for every u and v. The "only if" part follows from the property of k-connected graphs that any two k-subsets are linked. The "if" part follows from the observation that if $N_k(u)$ and $N_k(v)$ are linked then u and v are connected by k openly disjoint paths. Thus the linkedness subroutine needs be called at most $O(n^2)$ times.

2) But we do not even have to check the linkedness of $N_k(u)$ and $N_k(v)$ for every pair u, v of vertices, if we use the following simple lemma.

Lemma 4.1. *Let G be any graph and H, a k-connected graph with $V(H)=V(G)$. Then G is k-connected iff u and v are connected by k openly disjoint paths in G for every edge $uv\in E(H)$.* ▮

This implies that it suffices to check that $N_k(u)$ and $N_k(v)$ are linked for every edge $uv\in E(H)$. This means $O(nk)$ calls on a linkedness subroutine rather than $O(n^2)$.

3) Using Theorem 2.4, we shall be able to test whether a given k-tuple $N_k(u)$ is linked to *every* $N_k(v)$ faster than carrying out the corresponding linkedness test $n-1$ times. Calling this subroutine a "multilinkedness test for u", it is clear from the previous remarks that it suffices to check for the multilinkedness of k distinct vertices. This may be more efficient than than doing $O(nk)$ simple linkedness tests.

4) If we allow randomization then we can do even better. Let $r=r(n,k)$ be the least integer for which $\binom{n}{r}>n\binom{k-1}{r}$. It is easy to see that $r\leqq k$ and also $r\leqq$ $\leqq(n\log n)/(n-k)$. Now choose a set $Y\subset V(G)$, $|Y|=r(n,k)$ at random and then test for the multilinkedness of each $y\in Y$. If the test fails then the graph is of course not k-connected. If the test succeeds then it is still possible that the graph is not k-connected, but the set Y must then be contained in every cutset with fewer than k vertices. The probability that we made such an unfortunate choice for Y is at most $1/n$, by the definition of $r(n,k)$.

In the case of Theorem 2.4, the computation of the "physical" embedding in the proof requires solving a system of linear equations. For computational purposes it makes sense to solve the system in a finite field rather than in $\mathbf{R}$. Of course, this "modular" embedding has no physical or geometrical meaning any more, but the algebraic structure remains! If we apply Theorem 3.1 then it is quite natural to take finite fields right away.

Let us discuss the details of the algorithmic applications of Theorem 2.4 first. Consider a graph $G(V,E)$ with n nodes, and a set $X\subseteq V$. By a random (modular) X-embedding we mean the following. Choose at random a prime $p<n^5$. Choose uniformly at random a vector $c\in(\mathbf{Z}_p)^E$. Solve $(*)$ in $\mathbf{Z}_p$ to obtain $f=f_c$.

Lemma 4.2. *Let* $U\subset V$. *Then the probability that a random modular X-embedding f satisfies* $\operatorname{rank}(f(U))=p(U,X)$ *is at least* $1-n^{-3}$.

Proof. A standard application of Schwartz's Lemma (Schwartz 1980) and elementary number theory. Just note that the determinant D in the proof of Theorem 2.4 is a rational function of degree $2(n-k)k\leqq n^2$ in the coefficients in c, and its denominator never vanishes. ▮

Let $M(t)$ be the number of arithmetic steps required to multiply two $t\times t$ matrices. Recall that $M(t)=Q(t^{2.49\cdots})$ (Coppersmith and Winograd 1982). The following lemma is straightforward.

Lemma 4.3. (i) *Computing a random X-embedding f requires time* $O(M(n)\log d)$.

(ii) *For a subset* $U\subseteq V$, *computing* $\operatorname{rank}(f(U))$ *and finding a basis for the affine hull of* $f(U)$ *requires time* $O(M(|U|)\log d)$. ▮

From now on, assume we want to test whether or not a graph $G(V,E)$ on n vertices is k-connected. We may assume that the minimum degree in G is at least k. As above, choose an arbitrary k-subset $N_k(v)$ of $N(v)$, for every vertex $v\in V$. By the remarks above, we can state a version of Theorem 2.2 which is less appealing but more applicable in algorithms.

Theorem 4.4. *Let G be a graph and $k\geqq 0$, an integer. Then the following conditions are equivalent:*

(i) *G is k-connected.*

(ii) *For at least k distinct vertices $y \in V$, G has a convex $N_k(y)$-embedding f such that for every $v \in V \setminus N_k(y)$, $f(N_k(v))$ has full rank.* ▮

The following randomized algorithm takes a graph G and a vertex $y \in V$, computes a random $N_k(y)$-embedding and tests condition (ii) of Theorem 4.4.

Algorithm 4.5.

(1) Pick a random prime $p < n^5$ and a random vector $c \in (\mathbf{Z}_p)^{|E|}$.

(2) Compute a modular $N_k(y)$-embedding $f = f_c$ by solving the linear system $(*)$.

(3) For every $v \in V \setminus N_k(y)$ test if $\operatorname{rank}(f(N_k(v))) = k$; if satisfied for all such v, then return 'pass', else return 'fail'.

With large probability, this algorithm returns 'fail' if and only if there exists a $(k-1)$-element cut not containing y. So the only case in which it returns 'pass' for a non-k-connected graph is when y is contained in every cutset with fewer than k elements. By the discussion at the beginning of this section, it suffices to choose $r(n, k)$ distinct vertices y at random and run the above test for these vertices y. These considerations yield a Monte-Carlo algorithm to test k-connectivity.

Algorithm 4.6.

(1) Choose a set $Y \subset V$, $|Y| = r(n, k)$ at random.

(2) For each $y \in Y$, call the test in Algorithm 4.5 with the given G, y and k. If it returns 'fail' then print 'not k-connected' and halt.

(3) (If for all $y \in Y$, the test in Algorithm 4.5 returns 'pass' then) print 'k-connected'.

Theorem 4.7. *The complexity of Algorithm 4.6 is $O(n(\log n)^2 M(n-k)/(n-k) + {} + nM(k)\log n) = O(n^{2.5} + nk^{2.5})$. If G is k-connected [not k-connected], then A_k prints 'k-connected' ['not k-connected'] with probability larger then $1 - 1/n$.* ▮

Next we design a randomized algorithm that is somewhat less efficient, but never errs. Algorithm 4.6 may err in both directions. If all vertices y it tries happen to belong to every $(k-1)$-element cutset, then it may answer 'k-connected' when the graph is not. We mend this by making sure that at least one of the y's is not such a vertex, i.e. we select $|Y| = k$. On the other side, it may answer 'not k-connected', finding a degeneracy of some set $f(N_k(v))$ that is not due to a small cut but rather to bad random choices. We mend this by looking for a min-cut separating v from y, and, if not found, try new random choices.

To find a min-cut, we use the lattice structure of such cuts. For two subsets $X, U \subseteq V$, one can define a partial order among (X, U) min-cuts in which two cuts S_1, S_2 are related $(S_1 < S_2)$ if S_1 meets every path from S_2 to X. It is a well known fact (see [10]) that this partial order is a lattice. This lattice has a unique minimal element $S(X, U)$. The importance of this is that, although there may be many min-cuts separating U from X, $S(X, U)$ will determine the affine hull of $f(U)$. More exactly, we have the following lemma whose proof follows easily from this definition and the proof of Theorem 2.4.

Lemma 4.8. *Let f be a random X-embedding of the graph G. Let $U \subseteq V$ such that $p(U, X) = m < k$; let $S = S(U, X)$, and let T be the set of vertices separated from X by S (including the vertices of S). Let A be the affine hull of $f(U)$ Then $f(T) \subseteq A$ and with probability at least $1 - n^3/d$, $f(T) = A \cap f(V(G))$.*

The following refinement of Algorithm 4.5 takes a graph G, and a vertex $y \in V$, computes a random $N_k(y)$-embedding and tests condition (ii) of Theorem 4.4. In case of failure, it returns a set S of fewer than k vertices which separate y from some vertex v.

Algorithm 4.9.

(1) Pick a random prime $p < n^5$ and a random vector $c \in (\mathbf{Z}_p)^{|E|}$.

(2) Compute a modular $N_k(y)$-embedding $f = f_c$ by solving the linear system $(*)$.

(3) For every $v \in V \setminus N_k(y)$ test if rank $(f(N_k(v))) = k$; if satisfied for all such v, then return 'pass'.

(4) Else, if we find a $v \in V \setminus N_k(y)$ with rank $(f(N_k(v))) < k$, then compute $A = \mathrm{aff}\,(f(N_k(v)))$ and find the set S of those vertices which either belong to $N_k(y)$ or have a neighbor outside A;

(a) if $|S| < k$ then print 'not k-connected because of S' and halt;

(b) if $|S| \geqq k$ then start all over again with (1).

Using this algorithm as a subroutine, we obtain the following "Las Vegas" version of Algorithm 4.6.

Algorithm 4.10.

(1) Choose $Y \subseteq V$, $|Y| = k$ arbitrarily.

(2) For each $y \in Y$, call the test in Algorithm 4.9 with the given G, y and k. If it returns 'not k-connected because of S' then halt.

(3) (If for all $y \in Y$, the test in Algorithm 4.9 returns 'pass' then) print 'k-connected'.

The previous considerations contain the proof of the following theorem.

Theorem 4.11. *Algorithm* 4.10 *runs in expected time* $O((kM(n) + nkM(k)) \log n) = O(kn^{2.5} + nk^{3.5})$. *$G$ is k-connected iff it prints 'k-connected', i.e., it never errs.* ∎

By incorporating binary search it is easy to convert the Algorithms 4.6 and 4.10 into algorithms for finding the vertex connectivity $k(G)$ of the input graph G. We leave their specification to the reader, and only state that the Monte-Carlo algorithm has running time $O(n^{2.5} + nk(G)^{2.5})$, and errs with probability less that $1/n$. The Las Vegas algorithm has expected running time $O(k(G)n^{2.5} + nk(G)^{3.5})$.

Consider the problem of, for a given source $r \in V$, finding the number of (internally) disjoint paths to the vertex u, simultaneously for all $u \in V - \{r\}$.

Algorithm 4.12.

(1) Pick a random prime $p < n^5$ and a random vector $c \in (\mathbf{Z}_p)^{|E|}$, and compute an $N(r)$-embedding $f = f_c$;

(2) for each $u \in V - \{r\}$, print rank $(f(\overline{N}(u)))$.

Theorem 4.13. *Algorithm* 4.12 *runs in time* $O(n^{2.5} + n\Delta^{2.5})$ *where Δ is the maximum degree in G. It returns the correct answer with probability* $1 - 1/n$. ∎

Proof. The running time estimate is straightforward. To prove the error probability bound, let $p'(u, r)$ denote the number of internally disjoint (u, r)-paths, and let f be the $N(r)$-embedding. Then by Theorems 2.3 and 2.4,

(a) If $u \notin N(r)$, then $p'(u, r)=p(N(u), N(r))=\mathrm{rank}(f(N(u)))=\mathrm{rank}(f(\bar{N}(u)))$.

(b) If $u \in N(r)$, then $p'(u, r)=p(\bar{N}(u), N(r))=\mathrm{rank}\,(f(\bar{N}(u)))$. ∎

We could also use Theorem 3.1 to test a graph for connectivity by a Monte-Carlo algorithm. We choose a prime p and consider our matrices over $\mathbf{Z}_p$ (the prime need not be random for this application). We substitute random values for the x_{ij}. Then for any fixed $(n-k)\times(n-k)$ submatrix of A whose determinant is not identically 0, the probability that this submatrix will become singular by substitution is less than $1-n^2/p$ by Schwartz's Lemma. So we can check the k-connectivity between any two vertices by evaluating one $(n-k)\times(n-k)$ determinant. This leads to a Monte-Carlo algorithm whose running time is $O(n^2(\log n)^2 M(n-k)/(n-k))$, which, for large values of k, is better than Algorithm 4.6, but for small values of k, is much worse.

For small values of k, we can use the following trick: Compute the inverse D of the matrix A (after the random substitution). It is well known that an $(n-k)\times$ $\times(n-k)$ submatrix A' of A is non-singular iff the $k\times k$ submatrix of D formed by the rows and columns not occuring in A' is non-singular. Since the matrix D need be computed only once, this leads us to an $O(M(n) \log n+n(\log n)^2 M(k))$ Monte-Carlo algorithm, which is a very slight improvement upon Algorithm 4.6 for small k. We do not elaborate on these possibilities, however, in this paper.

These algorithms above lend themselves to parallelism. One can easily formulate analogous randomized parallel Monte-Carlo and Las Vegas algorithms, each running in time $O(\log^2 n)$ on an EREW PRAM ([14]), and use $nt(n)$ processors respectively, where $t(n)$ is the running time of the corresponding sequential algorithm. ($nM(n)$ is the best known bound on the number of processors needed to solve $n\times n$ linear systems over finite fields in parallel ([2]).)

Acknowledgement. We wish to thank Nick Pippenger and Éva Tardos for helpful discussions.

References

[1] A. V. Aho, J. E. Hopcroft and J. D. Ullman, *The Design and Analysis of Computer Algorithms*, Addison—Wesley, 1975.

[2] S. Berkowitz, On computing the determinant in small parallel time using a small number of processors, *Inform. Proc. Let.*, **18** (1984), 147—150.

[3] G. Birkhoff and S. MacLane, *A Survey of Modern Algebra*, MacMillan, 1970.

[4] D. Coppersmith and S. Winograd, On the asymptotic complexity of matrix multiplication, *SIAM J. Computing*, (1982), 472—492.

[5] R. Connelly, Rigidity and Energy, *Invent. Math.*, **66** (1982), 11—33.

[6] S. Even, *Graph Algorithms*, Computer Science Press, 1979.

[7] S. Even and R. E. Tarjan, Computing an *st*-numbering, *Theoret. Comp. Sci.*, **2** (1976), 339—344.

[8] Z. Galil, Finding the vertex connectivity of graphs, *SIAM J. Computing*, **9** (1980), 197—199

[9] A. W. Ingleton and M. J. Piff, Gammoids and transversal matroids, *J. Comb. Theory*, **B15**. (1973), 51—68.

[10] L. Lovász, Combinatorial Problems and Exercises, North-Holland, 1979.

[11] A. Lempel, S. Even and I. Cederbaum, An algorithm for planarity testing of graphs, *Theory of Graphs*, (1967), *International Symposium, Rome, P. Rosensfield ed.*, 215—232.

[12] H. Perfect, Symmetrized form of P. Hall's theorem on distinct representatives, *Qua Math. Oxford,* **17** (1966), 303—306.

[13] J. T. Schwartz, Fast probabilistic algorithms for verification of polynomial identiti *ACM* 27, **4** (1980), 701—717.

[14] V. Vishkin, Synchronous parallel computation, a survey, *TR* ♯ 71, *Department of Com Science, Courant Institute, NYU,* 1983.

[15] W. T. Tutte, How to draw a graph, *Proc. London Math. Soc.*, **13** (1963), 743—768.

N. Linial

The Hebrew University, Jerusalem and
Mathematical Sciences Research Institute,
Berkeley CA 94720, U.S.A.

L. Lovász

Eötvös Loránd University, H-1088 Budapest and
Mathematical Sciences Research Institute
Berkeley, CA 94720, U.S.A.

A. Wigderson

Mathematical Sciences Research Institute
Berkeley, CA 94720 and
The Hebrew University, Jerusalem, Israel