A PHYSICAL INTERPRETATION OF GRAPH CONNECTIVITY, AND ITS ALGORITHMIC APPLICATIONS

N. Linial
The Hebrew University, Jerusalem
Mathematical Sciences Research Institute, Berkeley

L. Lovász
Eötvös Lorand University, Budapest
Mathematical Sciences Research Institute, Berkeley

A. Wigderson
Mathematical Sciences Research Institute, Berkeley
The Hebrew University, Jerusalem

§0. Introduction

Connectivity is a basic property of graphs, and is related to other important concepts like reliability, communication and flow.

Connectivity is also one of the most well studied areas in graph theory. In this paper, we propose a new point of view on graph connectivity, based on geometric and physical intuition. Our main theorem is a geometric characterization of k-vertex connected graphs. It says that a graph G is k-connected if and only if G has a certain "nondegenerate convex embedding" in $\mathbb{R}^{k-1}$.

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

As an application of our theorem 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 $k^3n^{1.5}+k^2n^2$ [G]. Observe that our algorithms are faster for all $k \geq \sqrt{n}$, and for very dens grphs the Monte Carlo algorithm is faster by a linear factor!

Section 1 contains the main theorem and its proof. In Section 2 we describe the algorithmic applications. These include, in addition to the algorithms above, processor-efficient parallel counterparts.

Research supported in part by NSF Grant No. 8120790.

0272-5428/86/0000/0039$01.00 © 1986 IEEE

§1. Results

Notation.

Graph Theory

Let $G(V,E)$ be a graph. For a vertex $v \in V$, $N(v) = \{u \mid (v,u) \in E\}$ is 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!). G is k-connected if $|V| > k$ and for every two k-subsets X and Y, $p(X,Y) = k$. The largest k for which this holds is the vertex-connectivity of G, denoted $k(G)$. Other graph theoretic notation we use is standard.

Algebra

Let $\mathbb{R}^d$ denote the d-dimensional Euclidean space. Let $X = \{x_1,\ldots,x_m\}$ be a finite set of points in $\mathbb{R}^d$. The affine subspace of X is all points $\sum_{i=1}^{m} \lambda_i x_i$ with $\Sigma\lambda_i = 1$. The dimension of X, dim(X) is the dimension of this subspace over $\mathbb{R}$. X is in general position if $\dim(Y) = d$ for every $(d+1)$-subset $Y \subseteq X$. If X is not in general position, we call it degenerate. Let x_i have coordinates $x_{i1},\ldots,x_{id}$. Then rank(X) is the rank of the set of vectors $\{(1,x_{i1},x_{i2},\ldots,x_{id})\}_{i=1}^{m}$ in $\mathbb{R}^{d+1}$. We extend this definition to x_i with coordinates in any field F.

The point y is in the convex hull of X, denoted $y \in \mathrm{conv}(X)$, if $y = \sum_{i=1}^{m} \lambda_i x_i$ with some nonnegative λ_i such that $\Sigma_i \lambda_i = 1$.

Definition 1. Let $X \subseteq V$. A convex X-embedding of G is any mapping $f: V \longrightarrow \mathbb{R}^{|X|-1}$ s.t. for each $v \in V \setminus X$, $f(v) \in \mathrm{conv}(f(N(v)))$.

Theorem 1. The following two conditions on a graph G with $|V| > k > 1$ are equilvalent.

(1) G is k-connected.

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

Note that the special case $k = 2$ of our theorem is the well known s-t numbering of a 2-connected graph [E].

Theorem 1 will follow from Theorems 2 and 3 below.

Theorem 2. Let $X \subseteq V$. For every convex X-embedding f of G and every subset $U \subseteq V$, $U \neq \emptyset$, $\mathrm{rank}(f(U)) \leq p(U,X)$.

Theorem 3. There exists a convex X-embedding f, s.t. for every $U \subseteq V$, $U \neq \emptyset$, $\mathrm{rank}(f(U)) = p(U,X)$.

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

(1) $\Rightarrow$ (2) Assume G is k-connected and fix a k-subset $X \subseteq V$. Then Theorem 3 implies the existence of a convex X-embedding s.t. for every k-subset $Y \subseteq V$, $\text{rank}(f(Y)) = p(X,Y) \geq k$, so every k points are in general position.

Proof of Theorem 2. 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 $S \subseteq V$ s.t. $V \setminus S$ 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 $f(W) \subseteq \text{conv}(f(S))$. Note that this implies $\text{rank}(f(U)) \leq \text{rank}(f(W \cup S)) = \text{rank}(f(S)) \leq |S| = k = p(U,X)$.

To prove the claim, let $u \in W$. Hence $u \notin X$. Since f is a convex X-embedding, $f(u) \in \text{conv}(f(N(u))) \subseteq \text{conv}(f(W \cup S \setminus u))$, 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., $f(W) \subseteq \text{conv}(f(S))$.

Proof of Theorem 3. 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 springs. Glue the vertices of X to the extremes of a k-simplex in $\mathbb{R}^{k-1}$, and let the remaining vertices settle in a minimum energy equilibrium. It should be clear that if the potential carried by each spring 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 springs satisfy Hooke's Law) and our freedom in choosing the elasticity parameters (e.g. the thickness of the springs).

We proceed formally. Let $X = (x_0, x_1, \ldots, x_{k-1})$. Let e_0 be the zero vector and let e_i, $1 \leq i \leq k-1$ be the i-th unit vector in $\mathbb{R}^{k-1}$. An embedding $g: V \longrightarrow \mathbb{R}^{k-1}$ s.t. $g(x_i) = e_i$ for $0 \leq i \leq k-1$ is called an X-embedding. (Such an embedding is not necessarily convex!)

Assign to every edge $(u,v) \in E$ a positive elasticity coefficient $c_{uv} \in \mathbb{R}^+$, and let $c \in (\mathbb{R}^+)^E$ 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 (spring) $(u,v) \in E$ 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.

$$\sum_{u \in N(v)} c_{uv}(f(u)-f(v)) = 0 \text{ for all } v \in V \setminus X.$$

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., $f(x_i) = e_i$, $0 \leq i \leq k-1$), 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 $m \geq 1$. For a given vector c, f_c may not satisfy $\text{rank}(f_c(U)) = m$. 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 to be the following symmetric matrix (which was also used by Tutte [T])

$$T_{vv} = c_v \qquad v \in V$$

$$T_{uv} = \begin{cases} -c_{uv} & (u,v) \in E \\ 0 & (u,v) \notin E \end{cases}$$

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

$$T = \left[\begin{array}{c|c} T_1 & T_2^t \\ \hline T_2 & T_3 \end{array}\right]$$

where T_1 is the $k \times k$ matrix and T_3 is $(n-k) \times (n-k)$.

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 $f_c(x_i) = e_i$, $0 \leq i \leq k-1$ implies

$$F = \left[\begin{array}{c} 00\ldots0 \\ \hline I_{k-1} \\ \hline F_1 \end{array}\right]$$

and the stability condition $\frac{dP}{dg} = 0$ becomes $[T_2 \mid T_3]F = 0$, or equivalently

$$(*) \qquad T_3 F_1 = B, \qquad \text{where } B = -T_2 \begin{bmatrix} 0\ldots0 \\ I_{k-1} \end{bmatrix},$$

which 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 $M = \{u_1, u_2, \ldots, u_m\}$ and $p(M,X) = p(U,X) = m$. We want to show that $f_c(M)$ almost always has full rank. The volume of $f_c(M)$ in $\mathbb{R}^{k-1}$ is given by the determinant $D_c = (\det(AA^T))^{1/2}$, where the matrix A is

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

This determinant vanishes exactly when $f_c(M)$ does not have full rank. As the $f_i(u_j)$ 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 $P_1,P_2,\ldots,P_m$ from M to X, and assume without loss of generality that P_i connects u_i to x_i.

Let $c = (1,1,\ldots,1)$. If for this vector c, $D_c \neq 0$, then we are done. Else set $E' = \bigcup_{i=1}^{m} E(P_i)$. We claim that if we choose c_R s.t. the entries $c_e = 1$ for all $e \in E \setminus E'$, let $c_e = R$ for $e \in E'$, and let R tend to infinity, then the distances $\|f(u_i)-f(x_i)\|$ will tend to zero. Once they are small enough, the set $f_c(M)$ cannot have zero volume.

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 $f'(V(P_i)) = e_i$, $1 \leq i \leq m$ and (say) $f'(v) = e_0$ for any vertex not in X or $\cup P_i$. Then

$$P(f,c_R) \leq P(f',c_R) \leq 2\binom{n}{2} \leq n^2,$$

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

On the other hand,

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

By the Cauchy-Schwartz inequality and the fact that $|P_i| \leq n$ we get

$$\sum_{(u,v)\in E(P_i)} \|f(u)-f(v)\|^2 \geq \frac{1}{n}\Big\| \sum_{(u,v)\in E(P_i)} f(u)-f(v)\Big\|^2 = \frac{1}{2}\|f(u_i)-f(x_i)\|^2.$$

Let $\delta = \delta(R) = \max_{1\leq i\leq m} \|f(u_i)-f(x_i)\|^2$. Then $P(f,c_R) \geq \frac{R\delta}{n}$, and so $\delta \leq \frac{n^3}{R}$, and $\delta \to 0$ as $R \to \infty$.

A simple calculation shows that the determinant $D_{c_R} \neq 0$ whenever $\delta < \frac{1}{2}$, so it does not vanish for any c_R with $R > 2n^3$. Q.E.D.

§2. Algorithms

Our model of computation is the logarithmic cost RAM, see [AHU]. The computation of the "physical" embedding in the proof of Theorem 3 requires solving a system of linear equations. For computational purposes it makes sense to solve the system in a finite field rather than in $\mathbb{R}$. Of course, this "modular" embedding has no physical or geometrical meaning anymore, but the algebraic structure remains!

Fix G(V,E) with $|V| = n$, $X \subseteq V$, and d an integer. By a random (modular) X-embedding we mean the following. Choose at random a prime $p < d$. Choose uniformly at random a vector $c \in (\mathbb{Z}_p)^E$. Solve (*) in $\mathbb{Z}_p$ to obtain $f = f_c$.

Corollary 1. Fix G, $X,U \subseteq V_2$ and $d \in N$. Then for a random X-embedding f we have $\Pr[\ \text{rank}(f(U)) = p(U,X)\ ] \geq 1 - \frac{n}{d}$.

Proof. A standard application of Schwartz's Lemma [S] and the Chinese Remainder Theorem (e.g., [AHU]). Just note that the determinant D in the proof of Theorem 3 is a rational function of degree $2(n-k)k \leq n^2$ in the coefficients in c, and whose denominator never vanishes. For more details, see e.g., [RV].

Let M(t) be the number of arithmetic steps required to multiply two $t \times t$ matrices. Recall that $M(t) = t^{2.49\ldots}$ [CW]. The following lemma is straightforward.

Lemma 1. Let G, $X \subseteq V$ and $d \in N$ be given. Then
(i) Computing a random X-embedding f requires time $O(M(n)\log d)$.
(ii) For a subset $U \subseteq V$, computing rank(f(U)) and finding a basis for the affine hull of f(U) requires time $O(M(|U|)\log d)$.

Definition. A vertex $x \in V$ is critical if x belongs to every min-cut in G.

An immediate consequence of the definition is

Lemma 2. If G is not k-connected, and $y \in V$ is chosen at random, then

$$\Pr[\ y \text{ is not critical}\ ] > \frac{n-k}{n}.$$

From now on assume we want to test if a graph G(V,E) is k-connected, where $|V| = n > k$, and the minimum degree in G is at least k. For every vertex $v \in V$ let $N_k(v)$ denote an arbitrary k-subset of N(v). Now we can state a more "effective" version of Theorem 1.

Theorem 1'. The following conditions are equivalent:
(i) G is k-connected.
(ii) Let $y \in V$ be a non-critical vertex. G has a convex $N_k(y)$-embedding f s.t. for every $v \in V \setminus N_k(y)$, $f(N_k(v))$ has full rank. In fact, a random embedding will have this

property with probability $1 - \frac{n^3}{d}$.

The following randomized algorithm takes G, a vertex $y \in V$ and $d \in N$, computes a random $N_k(y)$-embedding and tests condition (ii) of Theorem 1'.

Algorithm $Test_k(G,y,d)$

(1) Pick a random prime $p < d$, and $c \in (\mathbb{Z}_p)^{|E|}$.

(2) Compute an $N_k(y)$-embedding $f = f_c$ using (*).

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

2.1. A Monte-Carlo Algorithm

We can use $Test_k$ and Theorem 1' to give a Monte-Carlo algorithm for testing k-connectivity.

Algorithm $A_k(G)$

(1) Repeat for $\frac{n \log n}{n-k}$ steps

  (a) choose $y \in V$ at random

  (b) if $Test_k(G,y,n^5)$ = 'fail' then print 'not k-connected' and halt.

(2) (if for all y, $Test_k(G,y,n^5)$ = 'pass' then) print 'k-connected'.

From the previous discussions we conclude:

Theorem 4.

(i) The complexity of A_k is $O(\frac{n(\log n)^2}{n-k}M(n-k) + nM(k)\log n) = O(n^{2.5+n \cdot k^{2.5}})$.

(ii) If G is k-connected, then A_k prints 'k-connected' with probability $> 1 - \frac{1}{n}$.

(iii) If G is not k-connected, then A_k prints 'not k-connected' with probability $> 1 - \frac{1}{n}$.

2.2. A Las-Vegas Algorithm

This subsection is devoted to a randomized algorithm that is somewhat less efficient, but never errs. Algorithm A_k may err in both directions. If all vertices y it tries are critical, it may answer 'k-connected' when the graph is not. We mend this by making sure one of the y's is not critical. 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. These ideas are captured in the following easy lemmas.

Lemma 3. Let $Y \subseteq V$ be any subset with $|Y| = k$. Then if G is not k-connected, at least one $y \in Y$ is not critical.

This trivially takes care of the first problem. 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 is closer to X than

S_2. It is a well known fact [L] that this partial order is a lattice. As a lattice has a unique minimal element, we have:

Lemma 4. Let $X,U \subseteq V$ s.t. $p(X,Y) = m < |X|$. Then there exists a unique cut set S, $|S| = m$ s.t.
(i) $G \setminus S$ has no (X,U) paths, and
(ii) if H is the subgraph of G with all components of $G \setminus S$ containing vertices from X, then for every $w \in V(H)$, $p(S \cup \{w\},X) = m+1$. (In other words, S is the min-cut which is closest to X.)

The importance of Lemma 4 is that, although there may be many min-cuts separating U from X, the one which is closest to X will determine the subspace in which f(U) lies. We use this in Lemma 5, which follows from Lemma 4, Theorem 3 and Corollary 1.

Lemma 5. Let f be a random X-embedding. Let $U \subseteq V$ s.t. $p(U,X) = m < k$; let $S = S(U,X)$ be the cut set from lemma 4, and let A be the affine subspace generated by f(U). Then
(i) If u is connected to U in $G \setminus S$, then $f(\bar{N}(u)) \subseteq A$.
(ii) For every $s \in S$, either $s \in X$ or with probability at least $1 - \frac{n^3}{d}$, $f(N(s)) \not\subseteq A$, i.e., s has a neighbor outside A.

Now we can put everything together.

Algorithm $B_k(G)$
(1) Choose $Y \subseteq V$, $|Y| = k$ arbitrarily.
(2) For each $y \in Y$ do;
  while $Test_k(G,y,n^5)$ = 'fail' do;
    let f be the embedding found in $Test_k$ and $v \in V$ the vertex s.t.
      $\text{rank}(f(N_k(v))) < k$.
    find a basis for A, the affine hull of $f(N_k(v))$.
    grow a maximal connected component Z of v s.t. for each $z \in Z$,
      $z \notin N_k(y)$ and $f(\bar{N}(z)) \subseteq A$.
    set $S = N(Z) \setminus Z$. If $|S| < k$ print 'not k-connected' and halt.
  end; (else if $|S| \geq k$ repeat the while loop)
end;
(3) (If for all $y \in Y$ some $Test(G,y,n^5)$ = 'pass' then) print 'k-connected'.

Theorem 5.
(i) Algorithm B_k runs in expected time $O((kM(n) + nkM(k))\log n) = O(kn^{2.5}+nk^{3.5})$.
(ii) G is k-connected iff B_k prints 'k-connected', i.e., B_k never errs.

Proof.
(i) If G is k-connected, then by Theorem 1' a random embedding will have no degeneracy

with probability $1 - \frac{1}{n}$ for every $y \in Y$. Similarly, by Lemma 5, if G is not k-connected, $1 - \frac{1}{n}$ is a lower bound on the probability that the set S found by the algorithm will indeed have cardinality less than k. Hence the time bound will be obeyed with probability $> 1 - n^{-b}$ for any fixed positive b.

(ii) If G is not k-connected, then B_k can certify its answer with a cut S, $|S| < k$. If G is k-connected, a convex embedding will certify it for at least one $y \in Y$.

2.3. More Algorithms

By incorporating binary search it is easy to convert the families of algorithms A_k and B_k into algorithms A and B (resp.) for finding the vertex connectivity of the input graph G. We leave their specification to the reader.

Theorem 6. Algorithm A is a Monte-Carlo algorithm s.t.

(i) A halts after $O(n^{2.5}+n \cdot k(G)^{2.5})$ steps.

(ii) If A outputs 'k', then $Pr[\ k(G)=k\ ] > 1 - \frac{1}{n}$.

Theorem 7. Algorithm B is a Las-Vegas algorithm that computes k(G) in expected time $O(k(G)n^{2.5} + n \cdot k(G)^{3.5})$.

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 C.

(1) Set $d = n^5$

(2) Compute a random N(r)-embedding f of G-{r}.

(3) Output $rank(f(\bar{N}(u)))$ for each $u \in V-\{r\}$.

Theorem 8.

(i) Algorithm C runs in time $O(n^{2.5}+n \cdot \Delta^{2.5})$ where Δ is the maximum degree in G.

(ii) Algorithm C returns the correct answer with probability $1 - \frac{1}{n}$.

Proof of (ii). Let p'(u,r) denote the number of internally disjoint (u,r)-paths, and let f be the N(r) embedding of Theorem 3.

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

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

The three algorithms above lend themselves to parallelism. One can easily formulate analogous randomized parallel algorithms A^p, B^p, C^p, each running in time $O(\log^2 n)$ on an EREW PRAM [V], and use $nt_A(n)$, $nt_B(n)$, $nt_C(n)$ processors respectively, where t_A, t_B, t_C are the running times of the related sequential algorithms. (nM(n) is the best known bound on the number of processors needed to solve n×n linear systems over finite fields in parallel [B].)

ACKNOWLEDGEMENT

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

REFERENCES

[AHU] A.V. Aho, J.E. Hopcropt and J.D. Ullman, The Design and Analysis of Computer Algorithms, Addison-Wesley, 1975.

[B] S. Berkowitz, "On Computing the Determinant in Small Parallel Time using a Small Number of Processors", IPL 18, pp.147-150, 1984.

[BM] G. Birkhoff and S. MacLane, A survey of Modern Algebra, MacMillan, 1970.

[CW] D. Coppersmith and S. Winograd, "On the Asymptotic Complexity of Matrix Multiplication", SIAM J. Computing, pp.472-492, 1982.

[E] S. Even, Graph Algorithms, Computer Science press, 1979.

[ET] S. Even and R.E. Tarjan, "Computing an st-numbering", TCS 2, pp.339-344, 1976.

[G] Z. Galil, "Finding the Vertex Connectivity of Graphs", SIAM J. Computing, 9, pp.197-199, 1980.

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

[LEC] A. Lempel, S. Even and I. Cederbaum, "An Algorithm for Planarity Testing of Graphs", Theory of Graphs, International Symposium, Rome, P. Rosensfield Ed., pp.215-232, 1967.

[S] J.T. Schwartz, "Fast Probabilistic Algorithms for Verification of Polynomial Identities", JACM 27, 4 pp.701-717, 1980.

[RV] M.O. Rabin and V.V. Vazirani, "Maximum Matchings in General Graphs through Randomization", submitted for publication.

[V] V. Vishkin, "Synchronous Parallel Computation, A Survey", TR#71, Department of Computer Science, Courant Institute, NYU, 1983.

[T] W.T. Tutte, "How to Draw a Graph", Proc. London Math. Soc., 3, 13, pp.743-768, 1963.