GRAPH COLORING USING EIGENVALUE DECOMPOSITION*

Bengt Aspvall†

John R. Gilbert††

TR 83-545
February 1983

Department of Computer Science
Cornell University
Ithaca, NY 14853

*A preliminary version of this report was presented at the SIAM Conference on Applied Linear Algebra in Raleigh, NC on April 26, 1982

†The work of this author was supported in part by National Science Foundation grant MCS-81-05556

††The work of this author was supported in part by National Science Foundation grant MCS-82-02948

Graph Coloring Using Eigenvalue Decomposition*

Bengt Aspvall†

John R. Gilbert‡

Computer Science Department

Cornell University

Ithaca, New York

February 1983

Abstract

Determining whether the vertices of a graph can be colored using k different colors so that no two adjacent vertices receive the same color is a well-known NP-complete problem. Graph coloring is also of practical interest (for example, in estimating sparse Jacobians and in scheduling), and many heuristic algorithms have been developed. We present a heuristic algorithm based on the eigenvalue decomposition of the adjacency matrix of a graph. Eigenvectors point out "bipartite-looking" subgraphs that are used to refine the coloring to a valid coloring. The algorithm optimally colors complete k-partite graphs and certain other classes of graphs with regular structure.

---

* A preliminary version of this report was presented at the SIAM Conference on Applied Linear Algebra in Raleigh, NC, on April 26, 1982.

† The work of this author was supported in part by National Science Foundation grant MCS-81-05556.

‡ The work of this author was supported in part by National Science Foundation grant MCS-82-02948.

I. Introduction

Discrete mathematics and combinatorial algorithms have made considerable contributions to numerical methods in recent years. Many of these contributions have come from graph theory; for example, graphs can be used to model sparse Gaussian elimination [Rose, 1972; George, 1977], and graph coloring gives efficient ways to estimate sparse Jacobian matrices [Coleman and Moré, 1983].

In this report we will turn this around and apply some numerical analysis to a problem in graph theory, namely coloring graphs. We will present two coloring heuristics based on the eigenvalues and eigenvectors of a graph's adjacency matrix. For the present, we do not claim that these heuristics are effective or efficient enough to compete with the various purely combinatorial coloring heuristics that exist. They do, however, offer a new view of the area where numerical and discrete computation overlap. Hence we believe that they are worth investigating further.

These heuristics are novel in that they use global information about the graph rather than local information; indeed, a small enough change in the graph does not change the behavior of the heuristics. We will discuss this point further in Section V. As might be expected, therefore, the heuristics perform best on graphs with very regular structures; Section IV goes into more detail.

The organization of the rest of this report is as follows. The next section reviews some necessary background in graph theory and linear algebra. Section III presents the basic ideas of the report and uses them in an algorithm to find an approximately correct two-coloring of a graph. Section IV presents an algorithm to find a correct coloring that may use more colors than necessary; it also describes some classes of graphs for which this algorithm finds a minimum coloring. The last section considers questions of stability: what happens if the graph changes slightly, and how accurate must the numerical calculations be to get the correct discrete answer? In this section we also discuss open problems and directions for further work.

II. Background

We begin with some standard definitions.

A *graph* $G = (V, E)$ consists of a set V of *vertices* and a set E of *edges*. An edge is an unordered pair $\{v, w\}$ of distinct vertices. If $\{v, w\}$ is an edge, vertices v and w are *adjacent*. Edge $\{v, w\}$ is *incident* on vertices v and w, which are its *endpoints*. The number of edges incident on a vertex is its *degree*. If all the vertices in a graph have the same degree d, the graph is *regular* of degree d.

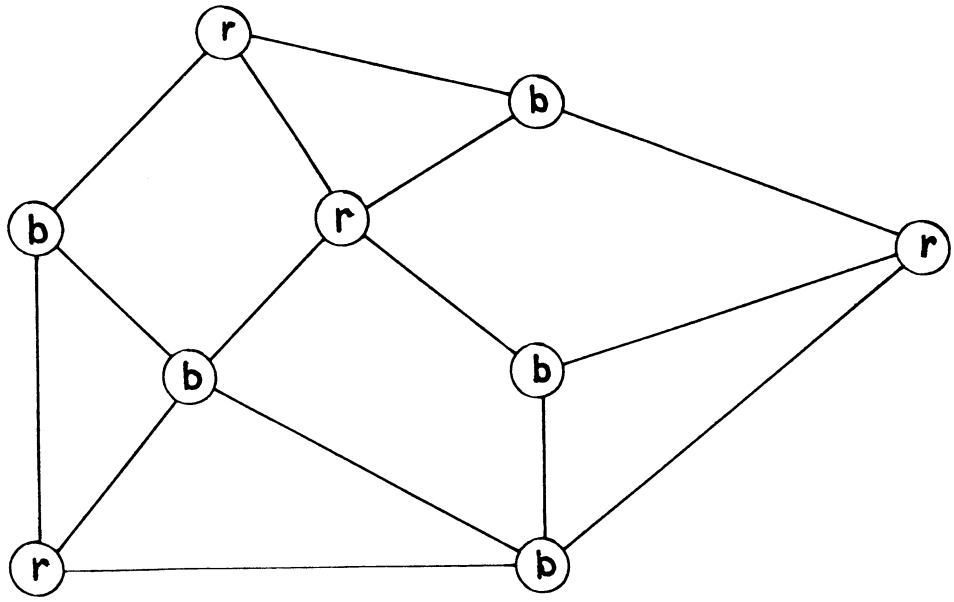

Figure 1. An approximate two-coloring.

A *path* of length k between vertices v and w is a sequence of vertices $v = v_0, v_1, \ldots, v_k = w$ such that $\{v_{i-1}, v_i\}$ is an edge for $1 \leq i \leq k$ and all the vertices $v_1, \ldots, v_k$ are distinct. If every pair of vertices in G is joined by a path, G is *connected*.

A *coloring* of a graph is an assignment of a color to each vertex. It is a *correct coloring* if every edge has different colors on its endpoints. We say that an edge *violates* or *satisfies* the coloring condition according to whether its endpoints have the same or different colors. A *minimum coloring* is a correct coloring with as few colors as possible; the number of colors in a minimum coloring of a graph G is its *chromatic number*, written $\chi(G)$.

If a minimum coloring can be found in polynomial time for every graph, then P = NP. We will be interested in polynomial time algorithms that find approximations to a minimum coloring. There are two kinds of approximations. An *approximate coloring* is a coloring that may not be correct; the fewer edges that violate the coloring condition, the better the approximation. In Figure 1, for example, eleven edges satisfy the coloring condition and four edges violate it. Any graph with m edges can be two-colored so that more than $m/2$ edges satisfy the coloring condition [Erdős and Kleitman, 1968], and such a coloring can be found in polynomial time. If the graph is regular, the coloring can be chosen so that half the vertices are of each color.

The other kind of approximation is an *approximately minimum correct coloring*, which is a correct

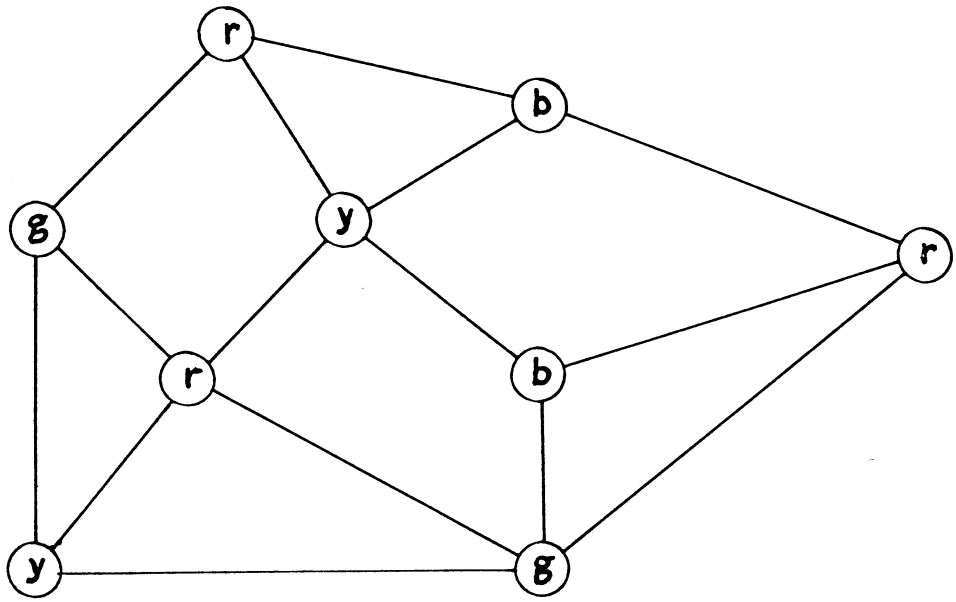

Figure 2. A correct but non-minimum coloring.

coloring; the fewer colors it uses, the better the approximation. In Figure 2, for example, we have used four colors to color a 3-colorable graph. Many heuristics of this sort have been proposed [Welsh and Powell, 1967; Matula, Marble, and Isaacson, 1972; Johnson, 1974; Brélaz, 1979]. Garey and Johnson [1976] showed that if a polynomial algorithm exists that always uses less than $2-\epsilon$ times the minimum number of colors, for any positive ϵ, then P = NP. Wigderson [1982] gave a polynomial algorithm that colors any n-vertex k-colorable graph with no more than

$$k^2 n^{(k-2)/(k-1)}$$

colors.

Another way to say that a graph has a correct k-coloring is that its vertices can be partitioned into k sets $V_1, \ldots, V_k$ such that no edge has both endpoints in the same set. Then the graph is called *k-partite*, the sets are called *parts*, and the graph is sometimes written $(V_1, V_2, \ldots, V_k, E)$. A *complete k-partite graph* is a k-partite graph in which every pair of vertices in different parts is joined by an edge.

If we number the vertices of graph G from 1 to n, the *adjacency matrix* of G is the n by n matrix $A = A(G)$ whose entry a_{ij} is 1 if vertices i and j are adjacent, 0 otherwise. Finding a correct k-coloring of G is equivalent to finding a permutation matrix P such that PAP^T has a k by k block structure in which the k diagonal blocks contain only zeroes.

The adjacency matrix is real, symmetric, and non-negative. Therefore it has n real *eigenvectors* $\mathbf{u}_1, \ldots, \mathbf{u}_n$. (We shall use italic uppercase for matrices, italic lowercase for scalars, and boldface lowercase for vectors. The components of the vector $\mathbf{x}$ are $x_1, \ldots, x_n$.) Furthermore, the eigenvectors can be chosen to be *orthonormal*, that is, so that the inner product $\mathbf{u}_i^T \mathbf{u}_j$ is zero if $i \neq j$ ($\mathbf{u}_i$ and $\mathbf{u}_j$ are *orthogonal*), and so that $\mathbf{u}_i^T \mathbf{u}_i = 1$ ($\mathbf{u}_i$ has *unit length*). The eigenvectors can also be chosen so the first component of each is nonnegative. With each eigenvector $\mathbf{u}_i$ is associated an *eigenvalue* λ_i such that $A\mathbf{u}_i = \lambda_i \mathbf{u}_i$. An eigenvalue is *simple* if it occurs only once. Since A is real and symmetric, its eigenvalues are real. The eigenvalues are the roots of the *characteristic equation* $\det(A - \lambda I) = 0$, which is a polynomial of degree n in λ.

The *spectral radius* $\rho(A)$ of A is $\max_i |\lambda_i|$, the largest magnitude of an eigenvalue. We shall number the eigenvalues and eigenvectors so that $\lambda_n \leq \lambda_{n-1} \leq \cdots \leq \lambda_1$. The sum of the diagonal elements of A is its *trace*, and is equal to the sum of the eigenvalues. Since the diagonal elements of an adjacency matrix are all zero, the sum of the eigenvalues is zero. Therefore $\lambda_1 \geq 0$ and $\lambda_n \leq 0$.

A matrix A is *reducible* if its rows and columns can be permuted symmetrically to place a block of zeroes in the lower left-hand corner, that is, if there is a permutation matrix P such that

$$PAP^T = \begin{pmatrix} C & D \\ 0 & E \end{pmatrix}.$$

The adjacency matrix of G is irreducible if and only if G is a connected graph with at least two vertices.

There is a rich theory of non-negative matrices; Varga [1962, Chapter 2] gives a good exposition. The results we need are due to Perron, Frobenius, and Gerschgorin, and we summarize them here along with the results mentioned above.

Theorem 1. Let G be a connected graph with $n > 1$ vertices and let A be G's adjacency matrix. Then

1. The eigenvalues $\lambda_1 \geq \cdots \geq \lambda_n$ of A are real.
2. The eigenvectors $\mathbf{u}_1, \ldots, \mathbf{u}_n$ can be chosen to be orthonormal.
3. $\sum_i \lambda_i = 0$.
4. $\lambda_1 \geq |\lambda_k|$ for all k; that is, λ_1 is the spectral radius of A.
5. $\lambda_1 > \lambda_2$; that is, the largest positive eigenvalue is simple.
6. $\mathbf{u}_1$ can be chosen to have positive components (we write $\mathbf{u}_1 > 0$).
7. λ_1 increases when any entry a_{ij} of A increases.
8. Either $\lambda_1 = \sum_j a_{ij}$ for all i or $\min_i \sum_j a_{ij} < \lambda_1 < \max_i \sum_j a_{ij}$.

All of these except (3) hold for any real, non-negative, symmetric, irreducible matrix.

Remark. Inequality (8) says that the spectral radius (and the largest positive eigenvalue) of G's adjacency matrix is bounded by the maximum and minimum degree of G's vertices. ∎

We can represent A in terms of its eigenvalues and eigenvectors as

$$A = \sum_k \lambda_k \mathbf{u}_k \mathbf{u}_k^T,$$

where the outer product $\mathbf{u}_k \mathbf{u}_k^T$ is an n by n matrix of rank one. If the sum is in decreasing order of $|\lambda_k|$, then the m-th partial sum $A^{(m)}$ minimizes $\|A - A^{(m)}\|_F$ over all rank-m matrices, where $\|B\|_F = \sqrt{\sum_{i,j} b_{ij}^2}$ is the Frobenius norm of the matrix B. In this case $\|A - A^{(m)}\|_F$ is equal to $\sqrt{\sum \lambda^2}$, where the sum is over the $n - m$ eigenvalues of smallest magnitude. Stewart [1973] gives details.

The *spectrum* of a graph is the multiset $\{\lambda_n, \ldots, \lambda_1\}$ of eigenvalues of its adjacency matrix. Cvetcović, Doob, and Sachs [1980] survey some known relationships between a graph's spectrum and chromatic number. Perhaps the most elegant is the following inequality, in which the lower bound is due to Hoffman [1970] and the upper bound to Wilf [1967].

Theorem 2. Let G be a graph with $n > 1$ vertices and let λ_1 and λ_n be its most positive and most negative eigenvalues. Then its chromatic number $\chi(G)$ satisfies

$$\frac{\lambda_1}{-\lambda_n} + 1 \leq \chi(G) \leq \lambda_1 + 1.$$

∎

The proof of this theorem is not constructive, and does not seem to lead to an efficient algorithm to color G with $\lambda_1 + 1$ colors.

Barnes and Hoffman [Barnes, 1982; Barnes and Hoffman, 1982] have used the eigenvalues and eigenvectors of a graph's adjacency matrix to partition the vertices into sets that have few edges between them. This is in a sense the dual of the coloring problem. Gould and other geographers [Straffin, 1980] have used the eigenvectors corresponding to positive eigenvalues to measure the "accessibility" of cities in trade networks; again, this is in a sense dual to the coloring problem.

Our investigation of spectral coloring heuristics was suggested by an algorithm due to Moler and Morrison [1983] that divides the letters of a cipher into vowels and consonants based on the singular values and singular vectors of the matrix of digram frequencies. They observed that, in English and several other languages, pairs of adjacent letters (digrams) are more likely to consist of a vowel and a consonant than two vowels or two consonants, so the weighted directed graph corresponding to the matrix of digram frequencies is approximately bipartite.

III. Two-colorings

The basic idea behind our coloring algorithm is best described in terms of bipartite graphs. In this section, we show how to color any bipartite graph correctly by examining the eigenvalue decomposition of its adjacency matrix. (Of course, bipartite graphs can be colored correctly in linear time using a simple depth-first search.) We then show how to use the ideas for approximate two-colorings.

Let $G = (V, E)$ be a bipartite graph with parts V_1 and V_2. By a suitable numbering of the vertices, the adjacency matrix $A(G)$ can be written in the form

$$A = \begin{pmatrix} 0 & B \\ B^T & 0 \end{pmatrix}.$$

Let $\mathbf{u}$ be an eigenvector of A with eigenvalue λ. If we write $\mathbf{u} = \begin{pmatrix} \mathbf{x} \\ \mathbf{y} \end{pmatrix}$, we have

$$\begin{pmatrix} 0 & B \\ B^T & 0 \end{pmatrix} \begin{pmatrix} \mathbf{x} \\ \mathbf{y} \end{pmatrix} = \lambda \begin{pmatrix} \mathbf{x} \\ \mathbf{y} \end{pmatrix}.$$

We claim that $\begin{pmatrix} \mathbf{x} \\ -\mathbf{y} \end{pmatrix}$ is also an eigenvector of A and its corresponding eigenvalue is $-\lambda$. Indeed, we have

$$\begin{pmatrix} 0 & B \\ B^T & 0 \end{pmatrix} \begin{pmatrix} \mathbf{x} \\ -\mathbf{y} \end{pmatrix} = \lambda \begin{pmatrix} -\mathbf{x} \\ \mathbf{y} \end{pmatrix} = -\lambda \begin{pmatrix} \mathbf{x} \\ -\mathbf{y} \end{pmatrix}.$$

Thus the spectrum of a bipartite graph is symmetric around zero. In fact, Sachs [Cvetković, Doob, and Sachs, 1980, Theorems 3.4, 3.11] proved a stronger result.

Theorem 3. The graph G is bipartite if and only if its eigenvalue spectrum is symmetric about the origin, and this happens if and only if $\lambda_1 = -\lambda_n$. ∎

From the Perron-Frobenius theorem, we know $\mathbf{u}_1 = \begin{pmatrix} \mathbf{x} \\ \mathbf{y} \end{pmatrix} > \mathbf{0}$. Thus the signs of the components of $\mathbf{u}_n = \begin{pmatrix} \mathbf{x} \\ -\mathbf{y} \end{pmatrix}$ correctly partition the vertices of G into the sets V_1 and V_2. We have the following algorithm for coloring bipartite graphs.

Algorithm 1 (*Two-coloring*). Color the vertices according to the signs of the components of the eigenvector $\mathbf{u}_n$ corresponding to the most negative eigenvalue λ_n. ∎

If the graph G is not bipartite, we can still partition its vertices by the signs of the components of $\mathbf{u}_n$. In the remainder of this section we give some intuition about why this might help find good colorings for arbitrary graphs. In the following section we will present an algorithm based on this idea and analyze its behavior on some classes of graphs.

Recall that the adjacency matrix A can be written in the form

$$A = \sum_k \lambda_k \mathbf{u}_k \mathbf{u}_k^T,$$

and that if the sum is in decreasing order of $|\lambda_k|$ then the m-th partial sum $A^{(m)}$ is the best rank-m approximation to A in the Frobenius norm. Now suppose that λ_n is the second largest eigenvalue in magnitude (i.e., $\lambda_1 > -\lambda_n \geq \lambda_2$), and let us take a look at the best rank-1 and rank-2 approximations to the adjacency matrix A. The matrix $\mathbf{u}_1\mathbf{u}_1^T$ has all positive elements, so $A^{(1)} = \lambda_1\mathbf{u}_1\mathbf{u}_1^T$ has all positive elements.

Since $\mathbf{u}_n$ is orthogonal to the positive vector $\mathbf{u}_1$, it must have both positive and negative components. Let $\binom{+}{-}$ denote a vector whose first components are positive and whose remaining ones are negative. (We will arbitrarily consider zero to be positive.) Let V_1 be the set of vertices corresponding to positive components and let V_2 correspond to negative components. By a suitable numbering of the vertices, we can write $\mathbf{u}_n = \binom{+}{-}$, so

$$\lambda_n\mathbf{u}_n\mathbf{u}_n^T = \lambda_n\begin{pmatrix} + & - \\ - & + \end{pmatrix} = \begin{pmatrix} - & + \\ + & - \end{pmatrix}.$$

When adding the second update $\lambda_n\mathbf{u}_n\mathbf{u}_n^T$ to $A^{(1)}$, obtaining the best rank-2 approximation $A^{(2)}$, we have

$$A^{(2)} = \lambda_1\mathbf{u}_1\mathbf{u}_1^T + \lambda_n\mathbf{u}_n\mathbf{u}_n^T = \begin{pmatrix} + & + \\ + & + \end{pmatrix} + \begin{pmatrix} - & + \\ + & - \end{pmatrix}. \tag{3.1}$$

The elements in the off-diagonal blocks in the two rank-1 matrices have the same signs and reenforce each other when added. In the diagonal blocks the elements added are of opposite signs and cancel each other. Thus, roughly speaking, we expect more edges between V_1 and V_2 than within V_1 and V_2, so we may view $A^{(2)}$ as the adjacency matrix for a "bipartite-looking" graph.

How good is this approximate two-coloring? Recall that every graph has an approximate two-coloring in which more than half the edges satisfy the coloring condition. The coloring above may not be this good. The following fairly weak result says that, for a class of graphs in which the most negative eigenvector partitions the vertices strongly enough, the approximate two-coloring cannot be too bad.

Theorem 4. Let G be connected and regular of degree $n/2$, and let its adjacency matrix A have eigenvalues $\lambda_n \leq \cdots \leq \lambda_1$ and eigenvectors $\mathbf{u}_n,\ldots,\mathbf{u}_1$. Let u_{jk} be the j-th component of eigenvector $\mathbf{u}_k$. If $|\lambda_n| \geq \lambda_2$ and there is a constant $0 < \eta < \sqrt{3}$ such that $|u_{in}/u_{jn}| < \eta$ for all i and j, then the number of edges that satisfy the coloring condition is more than $(3-\eta^2)/(2+2\eta^2)$ times the total number of edges.

Proof. By Theorem 1(8) the principal eigenvalue is $\lambda_1 = n/2$, and then $\mathbf{u}_1 = (1/\sqrt{n},\ldots,1/\sqrt{n})^T$. The best rank-1 approximation to A is $\lambda_1\mathbf{u}_1\mathbf{u}_1^T = \frac{1}{2}J$, where J is the matrix of all ones. Let

$$S = (s_{ij}) = 2(A - \lambda_1\mathbf{u}_1\mathbf{u}_1^T).$$

Then s_{ij} is 1 if vertices i and j are adjacent, -1 if not. Since λ_n is the eigenvalue of A of second largest magnitude, the best rank-1 approximation to S is $2\lambda_n\mathbf{u}_n\mathbf{u}_n^T$. Let $\mathbf{s} = \sqrt{-2\lambda_n}\,\mathbf{u}_n$. Then $2\lambda_n\mathbf{u}_n\mathbf{u}_n^T = -\mathbf{s}\mathbf{s}^T$, so $\mathbf{s}\mathbf{s}^T$ is the best rank-1 approximation to $-S$. That is, $\sum_{ij}(s_is_j + s_{ij})^2$ is minimum over all choices of $\mathbf{s}$.

Now

$$\sum_{ij}(s_i s_j + s_{ij})^2 = \sum_{ij} s_{ij}^2 + \sum_{ij} s_i^2 s_j^2 + 2\sum_{ij} s_i s_j s_{ij}. \tag{3.2}$$

The first sum on the right hand side above is n^2. The second is $\|\mathbf{s}\|^4$, which is $4\lambda_n^2$. Thus the third sum,

$$t = \sum_{ij} s_i s_j s_{ij},$$

is minimum over all choices of $\mathbf{s}$ with $\|\mathbf{s}\| = \sqrt{-2\lambda_n}$. In fact, we can evaluate t: the sum on the left in equation (3.2) is $4\|A - A^{(2)}\|_F^2$, which is $4\sum_{1\le i\le n}\lambda_i^2 - 4\lambda_1^2 - 4\lambda_n^2$. The latter sum is the square of the Frobenius norm of A, which is $n^2/2$. Since G is regular, λ_1 is $n/2$. Plugging this all into equation (3.2) gives $t = -4\lambda_n^2$. We will use only the fact that t is negative.

We divide the sum t into four parts. Let

$$\begin{aligned} a &= |\{(i,j) : s_i s_j < 0, s_{ij} = 1\}|, \\ b &= |\{(i,j) : s_i s_j > 0, s_{ij} = 1\}|, \\ c &= |\{(i,j) : s_i s_j < 0, s_{ij} = -1\}|, \\ \text{and } d &= |\{(i,j) : s_i s_j > 0, s_{ij} = -1\}|. \end{aligned}$$

Then a is twice the number of edges that satisfy the coloring condition, b is twice the number of edges that violate the coloring condition, c is twice the number of nonedges that satisfy the coloring condition, and d is twice the number of nonedges that violate the coloring condition. (Each edge and nonedge is counted once as (i,j) and once as (j,i), except that d counts each nonedge (i,i) just once.) Our goal is a lower bound on a. Since the graph is regular of degree $n/2$, we have

$$a + b = c + d = \frac{n^2}{2}. \tag{3.3}$$

Recall that the sum t above is negative. This sum has $b + c$ positive terms and $a + d$ negative terms. The ratio of the magnitudes of any two terms is less than η^2, so

$$b + c < \eta^2(a + d). \tag{3.4}$$

Now $\sum_i s_i = 0$ because $\mathbf{s}$ is parallel to $\mathbf{u}_n$ and hence orthogonal to $\mathbf{u}_1$. Therefore $\sum_{ij} s_i s_j = 0$. This sum has $b + d$ positive terms and $a + c$ negative terms. Thus

$$b + d < \eta^2(a + c). \tag{3.5}$$

Adding inequalities (3.4) and (3.5) and substituting $a+b$ for $c+d$ (by (3.3)) in the result yields $3b+a < \eta^2(3a + b)$. Rearranging terms gives

$$a > \frac{3 - \eta^2}{2 + 2\eta^2}(a + b). \tag{3.6}$$

Since a is twice the number of edges that satisfy the coloring condition and $a + b$ is twice the number of edges, this completes the proof. ∎

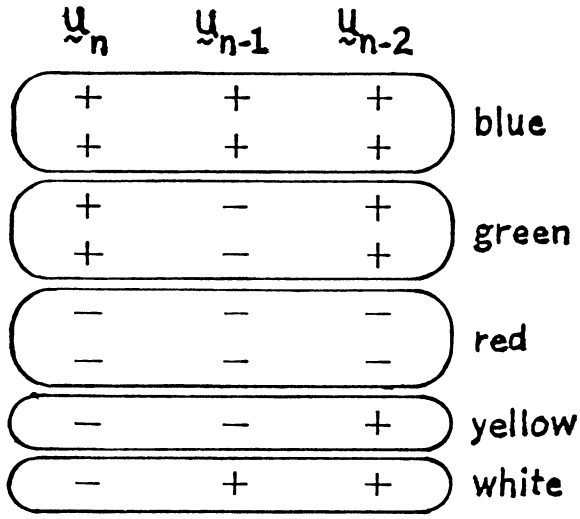

Figure 3. Partitioning vertices by sign pattern.

IV. A heuristic coloring algorithm

We can use the ideas of the last section in an algorithm to find an approximately minimum correct coloring of an arbitrary connected graph. The sign pattern of one eigenvector partitions the vertices into two sets. This gives an approximate two-coloring, which we refine to a valid coloring by partitioning the vertices according to additional eigenvectors.

Algorithm 2 (*Correct coloring*). Begin with all vertices the same color. Repeatedly select an eigenvector and use the signs of its components (with zero considered positive) to refine the coloring, until a correct coloring is obtained. ∎

For example, if the eigenvectors u_8, u_7, and u_6 of an 8-vertex graph have the sign patterns shown in Figure 3, the vertex partition they induce has 5 parts as shown.

The algorithm does not specify which eigenvectors to use. The discussion of low rank approximations above suggests that we select eigenvectors in increasing order of their eigenvalues, beginning with u_n. It turns out that, if we use enough eigenvectors, we eventually do get a correct coloring; indeed, we eventually color every vertex a different color.

Theorem 5. If the algorithm above is continued until it has used all the eigenvectors of $A = A(G)$, then every vertex is assigned a different color.

Proof. Consider the n by n matrix $U = (\mathbf{u}_1, \ldots, \mathbf{u}_n)$ whose columns are the eigenvectors of A. The color of vertex i is the sign pattern of row i of U. Since the columns of U are orthonormal, U is an orthogonal matrix and the rows $\mathbf{r}_1^T, \ldots, \mathbf{r}_n^T$ of U are also orthogonal; that is, $\mathbf{r}_i^T \mathbf{r}_j = 0$ if $i \neq j$. Eigenvector $\mathbf{u}_1$ has positive components, so the first components of $\mathbf{r}_i^T$ and $\mathbf{r}_j^T$ are both positive. Hence their inner product can be zero only if some other component is positive in one of them and negative in the other. Thus no two rows have the same sign pattern, so no two vertices have the same color. ∎

The discussion in the last section suggests that eigenvectors with negative eigenvalues should partition the vertices so that many edges satisfy the coloring condition, and eigenvectors with positive eigenvalues should partition the vertices so that few edges satisfy the coloring condition. That is, a negative eigenvalue divides the vertices into approximately independent sets, and a positive eigenvalue divides the vertices into approximately complete subgraphs. We conjecture that Algorithm 2 always finds a correct coloring after considering only the eigenvectors with negative eigenvalues.

Of course, a correct coloring that gives each vertex a different color is not very surprising. We now show that this algorithm finds minimum correct colorings for a class of graphs for which we know of no purely combinatorial polynomial-time coloring algorithm. To simplify the presentation, we first consider tripartite graphs.

Let G be a tripartite graph with parts V_1, V_2, and V_3. Let the parts have r, s, and t vertices respectively. Partition the adjacency matrix of G as

$$A = \begin{pmatrix} 0 & A_{12} & A_{13} \\ A_{12}^T & 0 & A_{23} \\ A_{13}^T & A_{23}^T & 0 \end{pmatrix}.$$

We call G ***block regular*** if in each block the row sum is constant and the column sum is constant. That is, $A_{ij}\mathbf{1} = b_{ij}\mathbf{1}$ and $A_{ij}^T\mathbf{1} = b_{ji}\mathbf{1}$, where $\mathbf{1}$ is the vector of all ones. In a block regular matrix, the number of edges between a given vertex v in V_i and vertices in V_j depends only on i and j (and this number is b_{ij}).

Theorem 6. Let G be a block regular tripartite graph. Then the adjacency matrix A has two eigenvectors with negative eigenvalues whose sign patterns correctly 3-color G.

Proof. The theorem follows from the following two lemmas. Let

$$B = \begin{pmatrix} 0 & b_{12} & b_{13} \\ b_{21} & 0 & b_{23} \\ b_{31} & b_{32} & 0 \end{pmatrix},$$

where b_{ij} is defined above. We call B the ***block degree*** matrix.

Lemma 1. Let A be the adjacency matrix of a block regular tripartite graph and let B be the corresponding block degree matrix. Then $(\alpha\mathbf{1}, \beta\mathbf{1}, \gamma\mathbf{1})^{\mathrm{T}}$ is an eigenvector of A with eigenvalue λ if and only if $(\alpha, \beta, \gamma)^{\mathrm{T}}$ is an eigenvector of B with eigenvalue λ.

Proof. Assume $(\alpha\mathbf{1}, \beta\mathbf{1}, \gamma\mathbf{1})^{\mathrm{T}}$ is an eigenvector of A with eigenvalue λ. We have

$$\begin{pmatrix} 0 & A_{12} & A_{13} \\ A_{12}^{\mathrm{T}} & 0 & A_{23} \\ A_{13}^{\mathrm{T}} & A_{23}^{\mathrm{T}} & 0 \end{pmatrix} \begin{pmatrix} \alpha\mathbf{1} \\ \beta\mathbf{1} \\ \gamma\mathbf{1} \end{pmatrix} = \begin{pmatrix} (\beta b_{12} + \gamma b_{13})\mathbf{1} \\ (\alpha b_{21} + \gamma b_{23})\mathbf{1} \\ (\alpha b_{31} + \beta b_{32})\mathbf{1} \end{pmatrix} = \lambda \begin{pmatrix} \alpha\mathbf{1} \\ \beta\mathbf{1} \\ \gamma\mathbf{1} \end{pmatrix}.$$

Looking at the right equality componentwise, we see that the first r equations are identical; so are the next s equations and the last t equations. Selecting one equation from each group, we have

$$\begin{pmatrix} 0 & b_{12} & b_{13} \\ b_{21} & 0 & b_{23} \\ b_{31} & b_{32} & 0 \end{pmatrix} \begin{pmatrix} \alpha \\ \beta \\ \gamma \end{pmatrix} = \lambda \begin{pmatrix} \alpha \\ \beta \\ \gamma \end{pmatrix}.$$

∎

Lemma 2. Let B be the block degree matrix for a block regular tripartite graph G. The sign patterns of two of the eigenvectors (with negative eigenvalues) of B partition the rows of B into singleton sets.

Proof. Since B is not necessarily symmetric, we cannot apply the Perron-Frobenius theorem directly to show that the eigenvectors partition the rows. However, we can use a similarity transform to get the desired result.

Counting the total number of edges between vertices in V_1 and V_2 in two ways, we see that $rb_{12} = sb_{21}$. Similarly, $rb_{13} = tb_{31}$ and $sb_{23} = tb_{32}$. Let

$$D = \begin{pmatrix} \sqrt{r} & 0 & 0 \\ 0 & \sqrt{s} & 0 \\ 0 & 0 & \sqrt{t} \end{pmatrix}.$$

Then

$$B' = DBD^{-1} = \begin{pmatrix} 0 & \sqrt{r/s}\, b_{12} & \sqrt{r/t}\, b_{13} \\ \sqrt{r/s}\, b_{12} & 0 & \sqrt{s/t}\, b_{23} \\ \sqrt{r/t}\, b_{13} & \sqrt{s/t}\, b_{23} & 0 \end{pmatrix}$$

is a symmetric nonnegative matrix with zero trace. Furthermore $D(\alpha, \beta, \gamma)^{\mathrm{T}}$ is an eigenvector of B' with eigenvalue λ if and only if $(\alpha, \beta, \gamma)^{\mathrm{T}}$ is an eigenvector of B with eigenvalue λ. Since D is diagonal with positive elements, the sign patterns of the eigenvectors of B are the same as the sign patterns of those of B'. Now the same argument as in the proof of Theorem 5 shows that these sign patterns partition the rows of B' (and those of B) into singletons. ∎

Theorem 6 says that the adjacency matrix A has two eigenvectors (with negative eigenvalues) whose sign patterns correctly partition the vertices, but it does not say which ones they are. We can find the right eigenvectors in polynomial time, unless one of B's negative eigenvalues λ has higher multiplicity as an eigenvalue of A than as an eigenvalue of B. In this case we may not be able to tell which of A's λ-eigenvectors correspond to B's λ-eigenvector(s).

Assuming no eigenvalue of B has higher multiplicity as an eigenvalue of A, however, we do not have to try all pairs of eigenvectors. We know that the negative sum of the corresponding eigenvalues is equal to the spectral radius $\rho(A)$. Thus we can find a minimum coloring of a block regular tripartite graph by trying at most half as many colorings as there are negative eigenvalues of its adjacency matrix.

The result generalizes straightforwardly to block regular k-partite graphs, but we may no longer be able to restrict our attention to eigenvectors with negative eigenvalues.

Theorem 7. Let G be a block regular k-partite graph. Then there is a set of at most $k-1$ eigenvectors whose sign patterns correctly partition the vertices of the graph. Furthermore, the negative sum of the corresponding eigenvalues is equal to the spectral radius $\rho(A(G))$. ∎

We now turn our attention to a restricted class of graphs for which we can show more precisely which eigenvectors partition the vertices correctly. Let G be a block regular k-partite graph for which b_{ij} depends only on j; that is, the off-diagonal elements in a column of the block degree matrix are all equal. We call such a graph ***strongly block regular.*** One example is a grid graph on a torus; see Figure 4.

Let the k parts have $n_1, n_2, \ldots, n_k$ vertices, where $n_1 \geq n_2 \geq \cdots \geq n_k$. Let $\alpha = b_{1k}$. Counting the number of edges between different partitions as we did in the proof of Lemma 2, we see that the block degree matrix B can be written as

$$B = \frac{\alpha}{n_k} \begin{pmatrix} 0 & n_2 & \ldots & n_k \\ n_1 & 0 & \ldots & n_k \\ \vdots & \vdots & & \vdots \\ n_1 & n_2 & \ldots & 0 \end{pmatrix}.$$

Except for the factor α/n_k, B looks like the block degree matrix B' for the complete k-partite graph with the same vertex partition. The eigenvectors of B and B' are thus the same, and their eigenvalues are related by the factor α/n_k. We therefore restrict our attention to complete k-partite graphs. We have the following theorem [Smith, 1970].

Theorem 8. A graph has exactly one positive eigenvalue if and only if its non-isolated vertices form a complete k-partite graph for some k. ∎

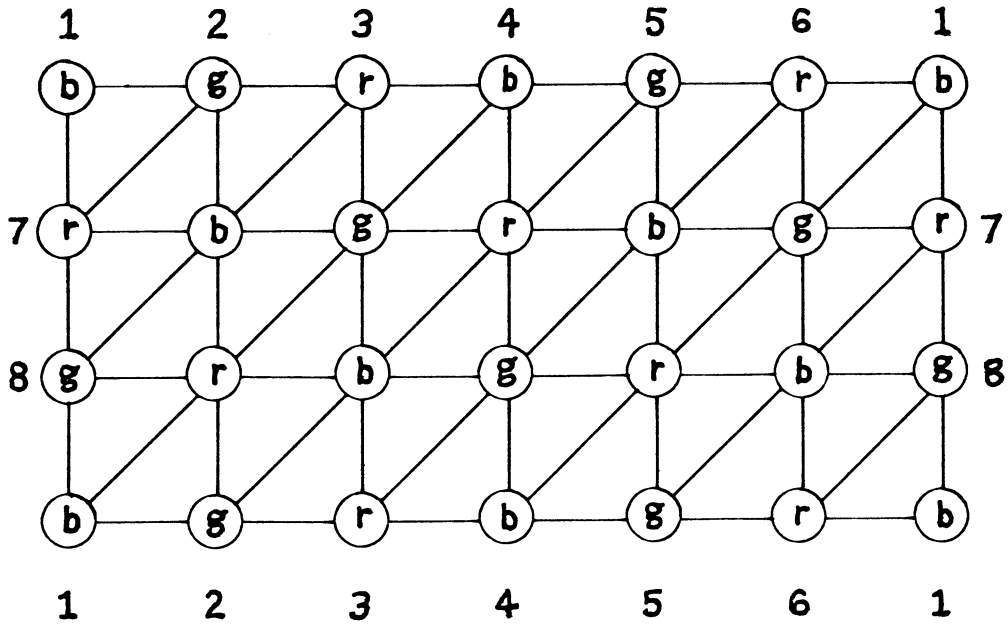

Figure 4. A strongly block regular graph (vertices with the same number are identical).

In fact, for complete k-partite graphs the characteristic equation is known [Cvetković, Doob, and Sachs, 1980, Section 2.6.8]:

$$p(\lambda) = \lambda^{n-k}\left(1 - \sum_{1\le i\le k} \frac{n_i}{\lambda + n_i}\right) \prod_{1\le j\le k} (\lambda + n_j).$$

From this equation, it follows that $p(-n_i) \times p(-n_{i+1}) \le 0$ with equality if and only if $n_i = n_{i+1}$. Thus the nonzero eigenvalues of a complete k-partite graph satisfy

$$-n_1 \le \lambda_n \le -n_2 \le \lambda_{n-1} \le \cdots \le \lambda_{n-k+2} \le -n_k < 0 < \lambda_1.$$

(Note that for strongly block regular graphs that are not complete k-partite graphs, the remaining $n-k-1$ eigenvalues of the adjacency matrix A might not be zero; in fact, they may be interleaved with the eigenvalues that are common with B.)

What can we say about the corresponding eigenvectors? Assume that $\lambda \ne -n_i$, for all $1 \le i \le k$. Let $\mathbf{u}$ be an eigenvector of the block degree matrix B' of a complete k-partite graph. Then $(B' - \lambda I)\mathbf{u} = \mathbf{0}$. By subtracting the first equation of $(B' - \lambda I)\mathbf{u} = \mathbf{0}$ from the j-th equation, it follows that $(n_1 + \lambda)u_1 = (n_j + \lambda)u_j$. Thus

$$\mathbf{u} = u_1(\lambda + n_1)\begin{pmatrix} 1/(\lambda + n_1) \\ 1/(\lambda + n_2) \\ \vdots \\ 1/(\lambda + n_k) \end{pmatrix}.$$

So for complete k-partite graphs the eigenvectors corresponding to negative eigenvalues have no zero components, and their blocks have the following sign patterns:

$$\mathbf{u}_n = (+,-,-,\ldots,-)^T, \mathbf{u}_{n-1} = (+,+,-,\ldots,-)^T, \ldots, \mathbf{u}_{n-k+2} = (+,+,+,\ldots,-)^T.$$

(If $n_i = n_{i+1}$ for some $1 \leq i < k$, then $\lambda_{n-i+1} = n_i$. In this case, there exists an eigenvector with zero components corresponding to vertices in certain block(s) in the partition.) From Theorem 7, we know that the sign patterns of $k-1$ of the eigenvectors correctly partition the vertices of any block regular graph. Therefore, for strongly block regular graphs, we now know exactly how the eigenvectors partion the vertices.

V. Conclusions

We have presented a new heuristic for coloring graphs. The approach is unusual in that it uses continuous mathematics for solving a combinatorial problem, and in this section we will discuss some of the implications. How accurately must the numerical computations be performed? Is our algorithm sensitive to perturbations in the input? That is, if the graph changes slightly, how does this affect the coloring? We will also sketch a way to improve a non-minimum coloring by "recycling" colors. Finally, we conclude with some open problems.

How accurately must the numerical computations be performed in finite precision arithmetic? That is, can we determine whether two eigenvalues λ and μ are equal, or whether a component of an eigenvector is positive, in time polynomial in the size of the graph? The eigenvalues are the zeroes of the characteristic polynomial $\det(A-\lambda I)$, which is a polynomial of degree n with coefficents bounded by $n!$ in magnitude. The eigenvalues are thus algebraic numbers of degree n. It follows from Theorem 1 of [Mignotte, 1982] that $\lambda \neq \mu$ implies $|\lambda - \mu| \geq \exp(-O(n^3 \log n))$. To test whether two eigenvalues are equal, we therefore need only to examine a polynomial number of bits. A similar argument holds for the components; that is, if $u_i \neq 0$, then u_i has at most a polynomial number of leading zeros. In theory, we can use any algorithm that is at least linearly convergent to compute the eigenvalue decomposition of a matrix. In practice, one might use the power method, for example, to obtain a few eigenvectors. For more details on computing the eigenvectors see [Stewart, 1973, Chapter 7; Parlett, 1980].

Our heuristic is based on global information about the graph in the sense that each eigenvector contains information about the entire adjacency matrix A. Changing the graph slightly will thus change all or almost all eigenvectors. The important point is that we expect the changes to be small. We present some experimental and some theoretical evidence for this.

First, we showed in Section IV that we obtain correct colorings of strongly block regular graphs. The strongly block regular graph in Figure 4 differs from a planar grid graph by only $O(\sqrt{n})$ edges. Our experiments show that various sorts of planar grid graphs are typically colored by Algorithm 2 in the minimum number of colors. Our experiments also indicate that Algorithm 1 gives, for many graphs, much better approximate two-colorings than Theorem 4 would lead us to believe. We plan to report on more extensive experiments in a future paper.

Second, the eigenvalue decomposition of a symmetric matrix is stable. That is, if a cluster of eigenvalues is well separated from the other eigenvalues, the subspace spanned by the eigenvectors corresponding to the cluster of eigenvalues is stable with respect to perturbation of A. (We refer the interested reader to [Davis and Kahan, 1970; Stewart, 1971; Stewart, 1973, Chapter 6] for a detailed discussion.) Thus, if the distinct eigenvalues of A are sufficiently well separated, then the subspace spanned by all the eigenvectors corresponding to an eigenvalue λ is insensitive to a perturbation. (If λ is a multiple eigenvalue, the eigenvectors corresponding to λ are not unique, but the subspace they span is unique.) Therefore, if a particular set of eigenvectors of $A(G)$ partitions the vertices of G correctly, we expect that the vertices of a slightly perturbed G' will be correctly or almost correctly partitioned by a corresponding set of eigenvectors of $A(G')$.

We know from Theorem 5 that Algorithm 2 will eventually find a correct coloring (although it might use n colors for some graphs). Suppose the algorithm returns a valid but non-minimum coloring. Then we can try to "recycle" colors as follows. Associate all vertices of the same color (they are all nonadjacent) with a ***supervertex***. Join two supervertices by an edge if there is a vertex associated with one of the supervertices that is adjacent to a vertex associated with the other supervertex. We then color the resulting graph. (If the resulting graph is small enough, an exhaustive algorithm might be feasible.) If two nonadjacent supervertices are colored with the same color, we can reduce the number of colors in the original coloring. This idea can be used recursively, and it does not depend on the particular coloring heuristic used.

As mentioned above, we conjecture that Algorithm 2 always colors a graph correctly (though not necessarily minimally) after considering only the eigenvectors with negative eigenvalues. If eigenvectors with negative eigenvalues partition the graph into pieces having few edges within them, then eigenvectors with positive eigenvalues can be viewed as partitioning the graph into pieces having few edges between them. (Notice what happens to Equation (3.1) when λ_n is replaced by $\lambda_2 > 0$.) Can this idea be used to find small separators in graphs, or perhaps to find large cliques?

In summary, we believe we have demonstrated that a numerical approach can sometimes give algorithms for purely combinatorial problems. Our main hope is to stimulate further research in the broad area of intersection between continuous and discrete mathematics.

References

Earl R. Barnes.
An algorithm for partitioning the nodes of a graph.
SIAM Journal on Algebraic and Discrete Methods **3**: 541–550, 1982.

Earl R. Barnes and Alan J. Hoffman.
Partitioning, spectra, and linear programming.
IBM Research Report RC 9511 (#42058), 1982.

Daniel Brélaz.
New methods to color the vertices of a graph.
Communications of the ACM **22**: 251–256, 1979.

Thomas F. Coleman and Jorge J. Moré.
Estimation of sparse Jacobian matrices and graph coloring problems.
SIAM Journal on Numerical Analysis, to appear, 1983.

Dragoš M. Cvetković, Michael Doob, and Horst Sachs.
Spectra of Graphs: Theory and Application.
Academic Press, 1980.

Chandler Davis and W. M. Kahan.
The rotation of eigenvectors by a permutation, III.
SIAM Journal on Numerical Analysis **7**: 1–46, 1970.

Paul Erdös and Daniel J. Kleitman.
On coloring graphs to maximize the proportion of multicolored k-edges.
Journal of Combinatorial Theory **5**: 164–169, 1968.

M. R. Garey and D. S. Johnson.
The complexity of near-optimal graph coloring.
Journal of the ACM **23**: 43–49, 1976.

J. Alan George.
Solution of linear systems of equations: Direct methods for finite element problems.
In V. A. Barker (editor), *Sparse Matrix Techniques: Copenhagen 1976*, Springer-Verlag Lecture Notes in Mathematics **572**: 52–101, 1977.

Alan J. Hoffman.
On eigenvalues and colorings of graphs.
In Bernard Harris (editor), *Graph Theory and Its Applications*, pages 79–91, Academic Press, 1970.

David S. Johnson.
Worst case behavior of graph coloring algorithms.
Proceedings of the 5th Southeastern Conference on Combinatorics, Graph Theory, and Computing, pages 513–527, Utilitas Mathematica, 1974.

David W. Matula, George Marble, and Joel D. Isaacson.
Graph coloring algorithms.
In Read [1972], pages 109–122.

Maurice Mignotte.
Identification of algebraic numbers.
Journal of Algorithms **3**: 197–204, 1982.

Cleve Moler and Donald Morrison.
Singular value analysis of cryptograms.
American Mathematical Monthly **90**: 78–87, 1983.

Beresford N. Parlett.
The Symmetric Eigenvalue Problem.
Prentice-Hall,1980.

Ronald C. Read, editor.
Graph Theory and Computing.
Academic Press, 1972.

Donald J. Rose.
A graph-theoretic study of the numerical solution of sparse positive definite systems of linear equations.
In Read [1972], pages 183–217.

John H. Smith.
Some properties of the spectrum of a graph.
In Richard Guy et al. (editors), *Combinatorial Structures and Their Applications*, pages 403–406, Gordon and Breach, 1970.

G. W. Stewart.
Error bounds for approximate invariant subspaces of closed linear operators.
SIAM Journal on Numerical Analysis **8**: 796–808, 1971.

G. W. Stewart.
Introduction to Matrix Computations.
Academic Press, 1973.

Philip D. Straffin, Jr.
Linear algebra in geography: Eigenvectors of networks.
Mathematics Magazine **53**: 269–276, 1980.

Richard S. Varga.
Matrix Iterative Analysis.
Prentice-Hall, 1962.

D. J. A. Welsh and M. B. Powell.
An upper bound for the chromatic number of a graph and its application to timetabling problems.
Computer Journal **10**: 85–86, 1967.

Avi Wigderson.
A new approximate graph coloring algorithm.
Proceedings of the Fourteenth Annual ACM Symposium on Theory of Computing, pages 325–329, 1982.

H. S. Wilf.
The eigenvalues of a graph and its chromatic number.
Journal of the London Mathematical Society **42**: 330–332, 1967.