SPECTRAL ANALYSIS OF LANDSCAPES*

PETTERI KASKI

1. Introduction

Fitness landscapes are mathematical structures of current interest in various disciplines ranging from theoretical biology to combinatorial optimization. Formally, a fitness landscape consists of three components [9]:

1. a set V of *configurations*;
2. a notion $\mathcal{W}$ of neighbourhood, nearness, distance, or accessibility on V; and
3. a *fitness function* $\psi : V \to \mathbb{R}$.

Typically we can assume that the pair $(V, \mathcal{W})$ has the structure of a finite undirected graph (which we shall call the *configuration graph* of the landscape), in which the configurations are the vertices, and the edges encode the symmetric "is a neighbour"-relation.

Under the previous assumption a fundamental observation is that we can view the fitness function ψ as a vector in $\mathbb{R}^{|V|}$, at which point the question arises whether it is possible to use the machinery of algebraic graph theory [5] to study ψ via the configuration graph.

Grover [7] first observed that many combinatorial optimization problems, when formulated as landscapes over a suitable configuration graph, have the property that the fitness function satisfies a difference equation which constrains its structure. In the language of algebraic graph theory this property translates to saying that the (zero-mean) fitness function, when viewed as a vector in $\mathbb{R}^{|V|}$, is an eigenvector of the configuration graph Laplacian. Such landscapes are called *elementary* [9].

In this manuscript we (1) illustrate the analysis of elementary landscapes by studying landscapes associated with the not-all-equal satisfiability problem (NAE-SAT) and the symmetric traveling salesman problem (symmetric TSP), both of which are NP-hard optimization problems [4]; and (2) survey some of the generic results that constrain the structure of elementary landscapes. These include Grover's results on the local minima and maxima of the fitness function and on the behaviour of greedy local search [7]; and the recent discrete nodal domain theorems [2], from which an upper bound on the number of nodal domains in an elementary landscape can be obtained.

The subsequent treatment is organized as follows. Section 2 contains the presupposed definitions and results. Section 3 studies two types of configuration graph: the n-dimensional cube, and the Cayley graph of the symmetric group generated by the set of all transpositions. These are the configuration graphs of the NAESAT and symmetric TSP landscapes, respectively, whose elementarity is established in Section 4. The final two sections review generic structural results on elementary

*Accompanying manuscript to the seminar presentation given in "T-79.300 Postgraduate Course in Theoretical Computer Science", Helsinki University of Technology, Spring term 2002.

landscapes. Section 5 presents Grover's results, and Section 6 outlines the discrete nodal domain theorems.

2. Preliminaries

2.1. Terminology for graphs. We briefly recall standard graph-theoretic terminology. (See [3, 5] for further reference.)

A (finite undirected) *graph* is a pair $X = (V, E)$, where V is a nonempty finite set of *vertices*, and E is a set of unordered pairs of distinct vertices called *edges*. We write $V(X)$ and $E(X)$ for the vertex and edge sets of a graph X, respectively.

The number of vertices in a graph is the *order* of the graph. Two vertices are *adjacent* if they are connected by an edge. A vertex is *isolated* if it is not adjacent to any other vertex. The number of vertices adjacent to a vertex is the *degree* of the vertex. We denote by $\deg(u)$ the degree of u. A graph is (k-)*regular* if all vertices have the same degree (k).

A graph $Y = (V', E')$ is a *subgraph* of a graph $X = (V, E)$ if $V' \subseteq V$, $E' \subseteq E$, and E' contains only edges with both endpoints in V'.

A *walk* in a graph is a nonempty sequence of vertices with the property that there is an edge between every pair of successive vertices in the sequence. The *length* of a walk is the length of the sequence minus one. A *path* in a graph is a walk with no repeated vertices in the sequence. The first and last vertices of a path are said to be connected by the path. A graph is *connected* if every pair of vertices is connected by a path. A *connected component* of a graph X is a connected subgraph that is not a proper subgraph of a connected subgraph of X. Thus, a graph is connected if and only if it has one connected component. The *diameter* of a connected graph is the maximum length of a shortest path connecting two vertices.

The *adjacency matrix* $A(X)$ of a graph X is the matrix with rows and columns indexed by the vertices of X such that the entry at row u, column v is defined by

$$A[u, v] = \begin{cases} 1 & \text{if } u \text{ and } v \text{ are adjacent; and} \\ 0 & \text{otherwise.} \end{cases}$$

In particular, the diagonal entries in $A(X)$ are 0.

The *incidence matrix* $B(X)$ of a graph X is the matrix with rows and columns indexed by $V(X)$ and $E(X)$, respectively, such that the entry at row u, column $\{v, w\}$ is defined by

$$B[u, \{v, w\}] = \begin{cases} 1 & \text{if } u \in \{v, w\}\text{; and} \\ 0 & \text{otherwise.} \end{cases}$$

If we direct the edges of a graph X by assigning for each edge $\{u, v\}$ either u or v as the start vertex, then we obtain a *directed incidence matrix* $D(X)$ from the incidence matrix $B(X)$ by reversing in each column the sign of the entry that corresponds to the start vertex.

2.2. Matrix theory. This section reviews the standard results from linear algebra that are required in what follows. (The results here are mostly extracted from Chapter 8 of [5]. For a general reference on linear algebra, see for example [6].)

The *characteristic polynomial* of a real square matrix A is defined by $p(\lambda) = \det(A - \lambda I)$. The zeros of the polynomial $p(\lambda)$ are the *eigenvalues* of A. The *(algebraic) multiplicity* of an eigenvalue λ is the multiplicity of λ as a zero of p.

An $n \times n$ real matrix has n eigenvalues, some of which may occur with multiplicity greater than one, and some of which may be complex.

A vector $x \neq 0$ that satisfies the equation $(A - \lambda I)x = 0$ is called an *eigenvector* of A. All eigenvectors associated with an eigenvalue λ together with the zero vector form a vector space, called the *eigenspace* of λ. The *(geometric) multiplicity* of an eigenvalue λ is the dimension of the corresponding eigenspace.

Theorem 2.1. *Let A be a real symmetric matrix. Then,*

(i) *the eigenvalues of A are real numbers; and*
(ii) *eigenvectors associated with distinct eigenvalues are orthogonal.*

Theorem 2.2. *Let A be a real symmetric $n \times n$ matrix. Then, $\mathbb{R}^n$ has an orthonormal basis consisting of eigenvectors of A.*

A real symmetric $n \times n$ matrix A is *positive semidefinite* if $u^T A u \geq 0$ for all $u \in \mathbb{R}^n$.

Theorem 2.3. *A real symmetric matrix A is positive semidefinite if and only if its eigenvalues are nonnegative.*

Theorem 2.4. *A real symmetric matrix A is positive semidefinite if and only if there exists a real matrix B such that $A = B^T B$.*

Theorem 2.5. *Let B be a real $m \times n$ matrix. Then,*

$$\dim \operatorname{Ker} B + \operatorname{rk} B = n.$$

Corollary 2.6. *Let B be a real $m \times n$ matrix. Then,*

(i) $\operatorname{Ker} B = \operatorname{Ker} B^T B$*; and*
(ii) $\operatorname{Ker} B^T = \operatorname{Ker} B B^T$*; and*
(iii) $\operatorname{rk} B = \operatorname{rk} B^T = \operatorname{rk} B^T B = \operatorname{rk} B B^T$.

The *spectrum* of a matrix A is a list of all of its eigenvalues together with their multiplicities. The *spectral radius* $\rho(A)$ of a matrix is the maximum of the moduli of its eigenvalues. The spectrum of a graph X is the spectrum of its adjacency matrix $A(X)$.

2.3. The Laplacian of a graph.

Definition 2.7. Let D be a directed incidence matrix of a graph X. The *Laplacian* of X is the matrix $L(X) = DD^T$.

The Laplacian is positive semidefinite by Theorem 2.4, and hence its eigenvalues are nonnegative by Theorem 2.3. Zero is always an eigenvalue of the Laplacian, and its multiplicity is the number of connected components in the graph as the corollary to the following theorem demonstrates.

Theorem 2.8. *Let X be a graph with n vertices and c connected components, and suppose that D is an directed incidence matrix of X. Then, $\operatorname{rk} D = n - c$.*

Proof. By Theorem 2.5 and Corollary 2.6 it suffices to prove that $\dim \operatorname{Ker} D^T = c$. Select any vector z such that $D^T z = 0$. By Definition of D we have $z_v - z_u = 0$ for every directed edge uv of X. Thus, z must be constant on every connected component of X. Since there are c connected components, we have $\dim \operatorname{Ker} D^T = c$. □

Corollary 2.9. *Let X be a graph with n vertices and c connected components. If L is the Laplacian of X, then* $\operatorname{rk} L = n - c$.

Proof. Let D be a directed incidence matrix of X. Theorem 2.8 gives $\operatorname{rk} D = n - c$, which combined with Corollary 2.6 gives $\operatorname{rk} L = \operatorname{rk} DD^T = \operatorname{rk} D = n - c$. □

The Laplacian is independent of the way the edges are directed as the following theorem demonstrates.

Theorem 2.10. *Let X be a graph, and denote by $\Delta(X)$ the diagonal matrix whose diagonal contains the degree of each vertex in X. Then,*

$$L(X) = \Delta(X) - A(X). \tag{1}$$

Proof. Let D be any directed incidence matrix of X. Then,

$$DD^T[u,v] = \sum_{\{s,t\}\in E(X)} D[u,\{s,t\}]D^T[\{s,t\},v] = \begin{cases} \deg(u) & \text{if } u = v\text{; and} \\ -1 & \text{if } \{u,v\} \in E(X)\text{; and} \\ 0 & \text{otherwise,} \end{cases}$$

for all vertices $u, v \in V(X)$, which establishes the claim. □

For k-regular graphs $\Delta(X) = kI$, so we obtain the following immediate corollary:

Corollary 2.11. *If X is a k-regular graph, then the adjacency matrix $A(X)$ and the Laplacian $L(X)$ have identical eigenspace structure. The eigenspace of $L(X)$ associated with eigenvalue λ_L corresponds to the eigenspace of $A(X)$ associated with eigenvalue λ_A, where*

$$\lambda_L = k - \lambda_A. \tag{2}$$

3. Configuration graphs

3.1. **The n-cube Q_2^n.** Denote by $\mathbb{Z}_2^n$ the set of vectors $x = (x_1, x_2, \ldots, x_n)$ of length n over $\{0,1\}$. The all-zero and all-one vectors are denoted by $\vec{0}$ and $\vec{1}$, respectively. The *Hamming distance* between two vectors $x, y \in \mathbb{Z}_2^n$ is the quantity

$$d_H(x,y) = |\{i \in \{1,\ldots,n\} \ : \ x_i \neq y_i\}|.$$

The *Hamming weight* of a word $x \in \mathbb{Z}_2^n$ is the quantity $w_H(x) = d_H(x, \vec{0})$.

Definition 3.1. The *n-cube* Q_2^n is the graph with vertex set $\mathbb{Z}_2^n$; two vertices $x, y \in \mathbb{Z}_2^n$ are connected by an edge if and only if $d_H(x,y) = 1$.

Example 3.2. The graphs Q_2^1, Q_2^2, Q_2^3, and Q_2^4 are depicted below.

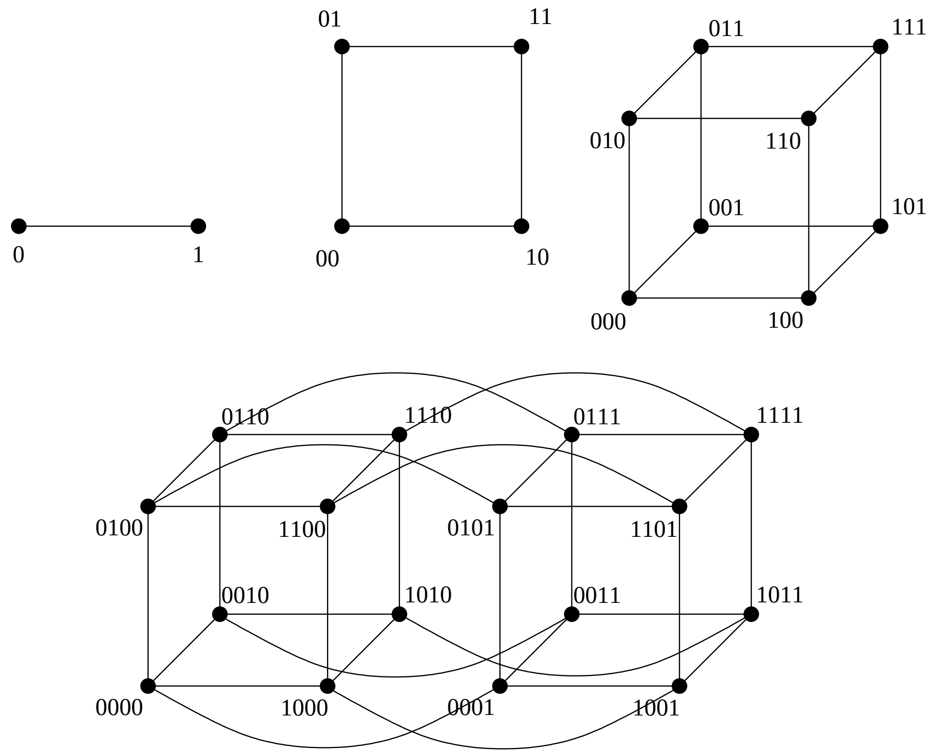

Theorem 3.3. *The n-cube Q_2^n is a connected bipartite n-regular graph with diameter n and order 2^n.*

Proof. The graph is bipartite since edges exist between vertices of even and odd Hamming weight only. The other properties are obvious. □

The spectrum of Q_2^n is well-known. Associate with each $z \in \mathbb{Z}_2^n$ a function $W_z : \mathbb{Z}_2^n \to \{-1, 1\}$ defined by the rule $x \mapsto (-1)^{\sum_{i=1}^n x_i z_i}$ for all $x \in \mathbb{Z}_2^n$. The functions $\{W_z \;:\; z \in \mathbb{Z}_2^n\}$ are known as the *Walsh functions.* The *weight* of a Walsh function W_z is $w_H(z)$.

Example 3.4. The Walsh functions for $n = 3$ are given in the table below.

x	000	001	010	011	100	101	110	111
W_{000}	1	1	1	1	1	1	1	1
W_{001}	1	−1	1	−1	1	−1	1	−1
W_{010}	1	1	−1	−1	1	1	−1	−1
W_{011}	1	−1	−1	1	1	−1	−1	1
W_{100}	1	1	1	1	−1	−1	−1	−1
W_{101}	1	−1	1	−1	−1	1	−1	1
W_{110}	1	1	−1	−1	−1	−1	1	1
W_{111}	1	−1	−1	1	−1	1	1	−1

Theorem 3.5. *The normalized Walsh functions $\tilde{W}_z = 2^{-n/2} W_z$ give an orthonormal basis of $\mathbb{R}^{2^n}$.*

Proof. Select any $x, z \in \mathbb{Z}_2^n$. Let $j_1, \ldots, j_{w_H(z)}$ be the indices of the 1-bits in z. In other words, $z_j = 1$ if and only if $j \in \{j_1, \ldots, j_{w_H(z)}\}$. Observe that $W_z(x)$ is the parity of the bits $x_{j_1}, \dots, x_{j_{w_H(z)}}$: an even number of 1-bits gives $W_z(x) = 1$, and an odd number of 1-bits gives $W_z(x) = -1$. Consequently,

$$\sum_{x \in \mathbb{Z}_2^n} W_z(x) = \begin{cases} 2^n & \text{if } z = \vec{0}\text{; and} \\ 0 & \text{otherwise,} \end{cases} \tag{3}$$

since for $z \neq \vec{0}$ there are equally many x with even and odd parities in $x_{j_1}, \dots, x_{j_{w_H(z)}}$. Now let $z, z' \in \mathbb{Z}_2^n$. By abuse of notation we denote by W_z also the column vector whose entries $W_z(x)$ are indexed by $x \in \mathbb{Z}_2^n$. We compute

$$W_z^T W_{z'} = \sum_{x \in \mathbb{Z}_2^n} W_z(x) W_{z'}(x) = \sum_{x \in \mathbb{Z}_2^n} (-1)^{\sum_{i=1}^n x_i(z_i + z_i')} = \sum_{x \in \mathbb{Z}_2^n} W_{z \oplus z'}(x), \tag{4}$$

where $\oplus$ denotes componentwise addition modulo 2. Combining (3) and (4), we obtain

$$W_z^T W_{z'} = \begin{cases} 2^n & \text{if } z = z'\text{; and} \\ 0 & \text{otherwise.} \end{cases}$$

Thus the vectors are mutually orthogonal, and hence linearly independent. We obtain a normalized basis by setting $\tilde{W}_z = 2^{-n/2} W_z$. □

Theorem 3.6. *Let A be the adjacency matrix of Q_2^n. Then, the distinct eigenvalues of A are $\lambda_k = n - 2k$, where $k = 0, 1, \dots, n$, and λ_k has multiplicity $\binom{n}{k}$. An orthonormal basis for the eigenspace of λ_k is given by the weight k normalized Walsh functions.*

Proof. Select any $z \in \mathbb{Z}_2^n$, and let $j_1, \dots, j_{w_H(z)}$ be the indices of the 1-bits in z. Fix a vertex $x \in \mathbb{Z}_2^n$ of Q_2^n. The vertices adjacent to x are precisely $y^{(1)}, \dots, y^{(n)}$, where

$$y_i^{(l)} = \begin{cases} x_i \oplus 1 & \text{if } i = l\text{; and} \\ x_i & \text{otherwise} \end{cases}$$

for all $i = 1, \dots, n$ and all $l = 1, \dots, n$. Thus,

$$W_z(y^{(l)}) = \begin{cases} W_z(x) & \text{if } l \notin \{j_1, \dots, j_{w_H(z)}\}\text{; and} \\ -W_z(x) & \text{if } l \in \{j_1, \dots, j_{w_H(z)}\}. \end{cases}$$

Summing over all l, we obtain

$$AW_z(x) = \sum_{l=1}^{n} W_z(y^{(l)}) = (n - w_H(z)) W_z(x) - w_H(z) W_z(x) = (n - 2w_H(z)) W_z(x).$$

Since x was arbitrary, $AW_z = (n - 2w_H(z)) W_z$. The claim now follows from Theorem 3.5 and the fact that there are $\binom{n}{k}$ binary vectors $z \in \mathbb{Z}_2^n$ with $w_H(z) = k$ for all $k = 0, \dots, n$. □

Example 3.7. The reader is invited to verify that the Walsh functions given in Example 3.4 are indeed the eigenfunctions of the graph Q_2^3 given in Example 3.2.

3.2. **The Cayley graph $\Gamma(S_n, T_n)$.** Recall that a *permutation* of a nonempty set E is a bijection of E onto E. The *symmetric group* S_n on $\{0, 1, \dots, n-1\}$ is the group formed by the set of all permutations of $\{0, 1, \dots, n-1\}$ with composition of permutations as the group operation. A *transposition* is a permutation that swaps two points and keeps the remaining points fixed. We denote the set of transpositions of $\{0, 1, \dots, n-1\}$ by T_n. The transposition that exchanges i and j, $i \neq j$, is denoted by $(i\ j)$.

Definition 3.8. The Cayley graph $\Gamma(S_n, T_n)$ is the graph with vertex set S_n; two vertices $\pi_1, \pi_2 \in S_n$ are connected by an edge if and only if $\pi_2^{-1}\pi_1 \in T_n$.

(The edge set is well-defined since $\pi_2^{-1}\pi_1 \in T_n$ if and only if $\pi_1^{-1}\pi_2 \in T_n$ as T_n is closed under taking of inverses.)

Example 3.9. The Cayley graphs $\Gamma(S_2, T_2)$, $\Gamma(S_3, T_3)$, and $\Gamma(S_4, T_4)$ are depicted below.

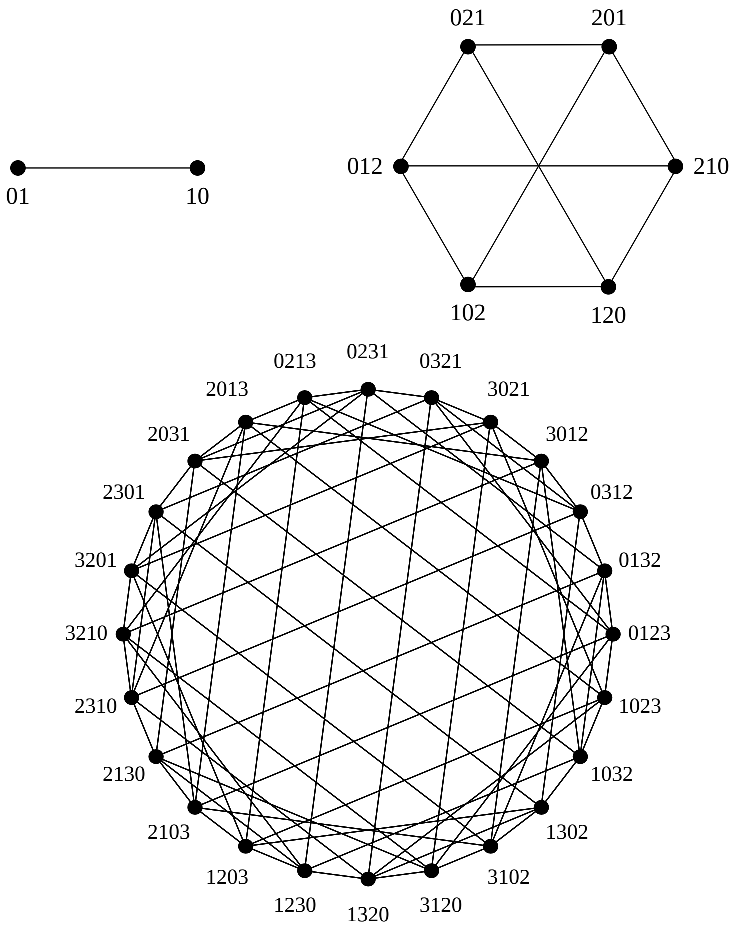

Theorem 3.10. *The Cayley graph $\Gamma(S_n, T_n)$ is a connected bipartite $\binom{n}{2}$-regular graph with diameter $n-1$ and order $n!$.*

Proof. Regularity is clear. The graph is connected since transpositions generate the symmetric group. The graph is bipartite since the endpoints of every edge consist of an even and an odd permutation. The diameter is at most $n-1$ since $n-1$ transpositions suffice to transform any permutation of $\{0, 1, \ldots, n-1\}$ into another. On the other hand, transforming for example the identity permutation to an n-cycle requires $n-1$ transpositions. □

The spectrum of $\Gamma(S_n, T_n)$ is known [1, 8] together with the associated eigenspace structure [10], however, its description requires a major excursion into the representation theory of finite groups, upon which we shall not embark in the present manuscript.

4. Two elementary landscapes

This section presents the elementarity proofs of Grover [7] for landscapes associated with NAESAT and symmetric TSP. (The proofs have been converted from their original form to the graph Laplacian formalism.)

4.1. **Not-all-equal satisfiability.** Let $V = \{v_1, \ldots, v_n\}$ be a set of n binary variables. A *literal* is either a variable $v \in V$ or its complement $\bar{v}$. A *clause* is a set of three literals. A *truth assignment* is a vector $x \in \mathbb{Z}_2^n$, which assigns the value x_i to the variable v_i for all $i = 1, \ldots, n$. A truth assignment *satisfies* a clause if and only if all of its literals do not have the same truth value, where it is assumed that the truth value of a complemented variable $\bar{v}$ is the opposite of the value assigned to v.

Definition 4.1. The decision problem Not-All-Equal Satisfiability (NAESAT) asks, given a set C of clauses over n binary variables, whether there exists a truth assignment to the variables that satisfies every clause in C.

In what follows we assume that the clauses which contain both a variable v and its complement $\bar{v}$ are removed from a NAESAT instance since they are obviously satisfied under any truth assignment.

A n-variable NAESAT instance defines in a natural way a landscape in which the configurations are the distinct truth value assignments, the configuration graph is the n-cube Q_2^n, and the cost of a truth value assignment $x \in \mathbb{Z}_2^n$ is the number of clauses **not** satisfied by x. We denote the cost of x by $\psi(x)$. Thus, $\psi(x) \geq 0$ with equality if and only if x satisfies all clauses in the instance.

The elementarity proof proceeds in two stages. First, we calculate the mean cost of a configuration (Theorem 4.2), and then establish that the zero-mean cost function is an eigenfunction of the Laplacian of Q_2^n (Theorem 4.3).

Theorem 4.2.

$$\frac{1}{2^n} \sum_{x \in \mathbb{Z}_2^n} \psi(x) = \frac{|C|}{4}.$$

Proof. Denote by $\psi_{C_i}(x)$ the cost of a single clause C_i from C. In other words,

$$\psi_{C_i}(x) = \begin{cases} 0 & \text{if } x \text{ satisfies } C_i\text{; and} \\ 1 & \text{otherwise.} \end{cases}$$

Since exactly 2 of the 8 possible truth value assignments to the three literals in a clause satisfy it, we have

$$\sum_{x\in\mathbb{Z}_2^n}\psi(x)=\sum_{i=1}^{|C|}\sum_{x\in\mathbb{Z}_2^n}\psi_{C_i}(x)=\sum_{i=1}^{|C|}2^{n-3}\cdot 2=|C|2^{n-2}.$$

Divide both sides by 2^n to obtain the claim. □

Now define the *zero-mean cost function* by

$$\bar{\psi}(x)=\psi(x)-|C|/4.$$

Theorem 4.3. *Let L be the Laplacian of the n-cube Q_2^n. Then,*

$$L\bar{\psi}(x)=4\bar{\psi}(x).$$

for all $x\in\mathbb{Z}_2^n$.

Proof. Denote again by $\psi_{C_i}(x)$ the cost of a single clause C_i from C. From (1) and the definition of Q_2^n we obtain for every $x\in\mathbb{Z}_2^n$:

$$L\bar{\psi}(x)=\sum_{d_H(x,y)=1}\big(\bar{\psi}(x)-\bar{\psi}(y)\big)=\sum_{i=1}^{|C|}\sum_{d_H(x,y)=1}\big(\psi_{C_i}(x)-\psi_{C_i}(y)\big). \tag{5}$$

Now,

$$\sum_{d_H(x,y)=1}\big(\psi_{C_i}(x)-\psi_{C_i}(y)\big)=\begin{cases}-1 & \text{if } x \text{ satisfies } C_i\text{; and}\\ 3 & \text{otherwise,}\end{cases}$$

or more compactly

$$\sum_{d_H(x,y)=1}\big(\psi_{C_i}(x)-\psi_{C_i}(y)\big)=4\psi_{C_i}(x)-1. \tag{6}$$

Substituting (6) into (5), we obtain

$$L\bar{\psi}(x)=4\psi(x)-|C|=4\bar{\psi}(x).$$

□

Thus, by Theorem 3.6 and Corollary 2.11, $\bar{\psi}$ is a vector in the eigenspace of L spanned by the weight 2 Walsh functions.

Corollary 4.4. *For any NAESAT instance over n variables, the cost function ψ can be written in the form*

$$\psi(x)=\frac{|C|}{4}W_{\vec{0}}(x)+\sum_{w_H(y)=2}\alpha_y W_y(x),$$

where $\alpha_y\in\mathbb{R}$.

Computing the coefficients α_y is straightforward because the coefficients for a clause set C are obtained as a sum of the coefficients for the individual clauses, i.e., the transformation from clauses to coefficients is linear.

4.2. **The traveling salesman problem.** Let A be any $n \times n$ nonnegative integer matrix whose rows and columns are indexed by n cities, which we shall label $\{0, 1, \ldots, n-1\}$. The entry $a_{i,j}$ of A at row i, column j is the *distance* from city i to city j. The distances are assumed to be symmetric, that is, $a_{i,j} = a_{j,i}$ for all i, j.

A *tour* is a permutation $\pi \in S_n$ of the n cities. The *length* of a tour π is

$$l(\pi) = a_{\pi(0),\pi(1)} + a_{\pi(1),\pi(2)} + \cdots + a_{\pi(n-2),\pi(n-1)} + a_{\pi(n-1),\pi(0)}.$$

Definition 4.5. Given a symmetric nonnegative intercity distance matrix A, the SYMMETRIC TRAVELING SALESMAN PROBLEM (symmetric TSP) asks for the length of the shortest tour (that visits all cities exactly once and finally returns to the initial city).

One possible way to associate a landscape with the symmetric TSP problem is to consider the tours as the distinct configurations, and regard two tours as adjacent if one can be obtained from another by transposition of two cities. The induced configuration graph in this case is precisely the Cayley graph $\Gamma(S_n, T_n)$.

The natural cost of a tour π is its length, that is, we define the cost function ψ by setting $\psi(\pi) = l(\pi)$ for all $\pi \in S_n$.

We next show that the resulting landscape is elementary, and start by computing the mean cost of a configuration:

Theorem 4.6.

$$\frac{1}{n!} \sum_{\pi \in S_n} \psi(\pi) = \frac{2}{n-1} \sum_{i<j} a_{i,j}. \tag{7}$$

Proof. Fix any i, j such that $0 \leq i < j \leq n-1$. There are exactly $n(n-2)!$ permutations $\pi \in S_n$ for which $\pi(k) = i$ and $\pi(k+1) = j$ hold for some $0 \leq k \leq n-1$. (The addition $k+1$ is performed modulo n.) By symmetry, both $a_{i,j}$ and $a_{j,i}$ appear exactly $n(n-2)!$ times in the sum on the right hand side of (7). Thus,

$$\frac{1}{n!} \sum_{\pi \in S_n} \psi(\pi) = \frac{1}{n!} \sum_{i<j} n(n-2)!(a_{i,j} + a_{j,i}) = \frac{2}{n-1} \sum_{i<j} a_{i,j}.$$

□

The zero-mean cost function is thus defined by

$$\bar{\psi}(\pi) = \psi(\pi) - \frac{2}{n-1} \sum_{i<j} a_{i,j}. \tag{8}$$

for all $\pi \in S_n$.

Theorem 4.7. *Let L be the Laplacian of the Cayley graph $\Gamma(S_n, T_n)$. Then,*

$$L\bar{\psi}(\pi) = 2(n-1)\bar{\psi}(\pi).$$

for all $\pi \in S_n$.

Proof. Fix any tour $\pi \in S_n$. By definition of the Laplacian

$$L\bar{\psi}(\pi) = \sum_{i<j} \bar{\psi}(\pi) - \bar{\psi}(\pi(i\ j)).$$

Select any i, j such that $0 \leq i, j \leq n-1$, and put $\delta_{ij} = \bar{\psi}(\pi) - \bar{\psi}(\pi(i\ j))$. For convenience we denote $a_{\pi(i),\pi(j)}$ simply by $l_{i,j}$. Moreover, all arithmetic on the indices

i, j is implicitly assumed to be modulo n to keep the notation simple. Without loss of generality we may take $l_{i,i} = 0$ for all i. When $j - i \not\equiv \pm 1$, we have (Case A on the figure below)

$$\delta_{ij} = (l_{i-1,i} + l_{i,i+1} + l_{j-1,j} + l_{j,j+1}) - (l_{i-1,j} + l_{j,i+1} + l_{j-1,i} + l_{i,j+1}).$$

When $j - i \equiv \pm 1$, we have (Case B on the figure below)

$$\delta_{ij} = (l_{i-1,i} + l_{i,i+1} + l_{j,j+1}) - (l_{i-1,j} + l_{j,i} + l_{i,j+1}).$$

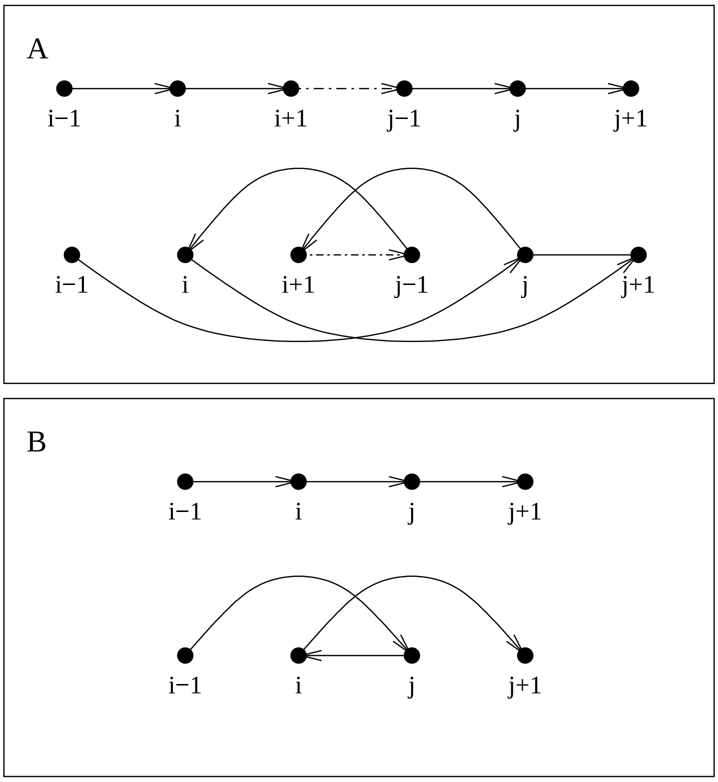

Thus, an expression valid for all i, j (including $i \equiv j$) is

$$\begin{aligned}\delta_{ij} =& (l_{i-1,i} + l_{i,i+1} + l_{j-1,j} + l_{j,j+1}) - (l_{i-1,j} + l_{j,i+1} + l_{j-1,i} + l_{i,j+1}) \\ & - (\Delta_{j,i+1} + \Delta_{j,i-1})(l_{j,i} + l_{i,j}),\end{aligned}$$

where $\Delta_{k,l} = 1$ if $k \equiv l$, and $\Delta_{k,l} = 0$ otherwise. Summing over all $0 \le i < j \le n-1$, it follows that

$$\begin{aligned} L\bar{\psi}(\pi) &= \sum_{i<j} \delta_{ij} = \frac{1}{2}\sum_{i,j} \delta_{ij} = \frac{1}{2}\Big(4n\psi(\pi) - 4\sum_{i,j} l_{i,j} - 4\psi(x)\Big) \\ &= 2(n-1)\psi(x) - 4\sum_{i<j} l_{i,j} = 2(n-1)\bar{\psi}(\pi), \end{aligned}$$

where the last equality is obtained using (8). □

5. Grover's results on elementary landscapes

Throughout this section we assume that (X, ψ) is an elementary landscape, where $X = (V, E)$ is the configuration graph, and ψ is the cost function, whose zero-mean counterpart is denoted by $\bar{\psi}$. L denotes the Laplacian of X, and $\lambda > 0$ is the eigenvalue of L associated with $\bar{\psi}$, that is, $L\bar{\psi} = \lambda\bar{\psi}$. (Note that we assume $\lambda > 0$; the case $\lambda = 0$ is omitted from consideration since $L\bar{\psi} = 0$ implies that $\bar{\psi}$ is constant on each connected component of X.)

5.1. **Local minima and maxima.** Elementary landscapes have the interesting property that the local minima and maxima are constrained to be below and above the average cost, respectively.

Definition 5.1. A vertex $u \in V$ is a *local cost minimum* if $\psi(u) < \psi(v)$ holds for all $\{u, v\} \in E$. Similarly, $u \in V$ is a *local cost maximum* if $\psi(u) > \psi(v)$ holds for all $\{u, v\} \in E$.

Theorem 5.2. *If $u \in V$ is a local cost minimum, then $\bar{\psi}(u) \leq 0$. If $u \in V$ is a local cost maximum, then $\bar{\psi}(u) \geq 0$. Equality holds if and only if u is isolated.*

Proof. We consider only the case of a local cost minimum; the proof for a local cost maximum is similar. If u is a local cost minimum, $\psi(u) - \psi(v) = \bar{\psi}(u) - \bar{\psi}(v) < 0$ for all $\{u, v\} \in E$. (Note that if u is isolated, i.e., has no adjacent vertices, then u is always a local minimum and maximum. Elementarity forces then $\bar{\psi}(u) = 0$.) By definition of the Laplacian and the fact that (X, ψ) is elementary,

$$\sum_{\{u,v\} \in E} (\bar{\psi}(u) - \bar{\psi}(v)) = L\bar{\psi}(u) = \lambda\bar{\psi}(u).$$

So, since $\lambda > 0$, we have that $\bar{\psi}(u) \leq 0$. Equality holds if and only if $\deg(u) = 0$. □

5.2. **Behaviour of greedy local search.** The following theorem and subsequent discussion demonstrate that greedy local search on an elementary landscape will relatively quickly reach an average cost configuration, no matter on what configuration the search is started.

We discuss only greedy local search for the minimum; local search for the maximum is handled similarly.

The following claim demonstrates that a nonisolated vertex with positive cost always has a neighbour with lesser cost. (So, greedy search for the minimum will by definition proceed to a vertex with at most that cost.)

Theorem 5.3. *For every nonisolated vertex $u \in V$, there exists a vertex $v \in V$ adjacent to u such that*

$$\bar{\psi}(v) \leq \left(1 - \frac{\lambda}{\deg(u)}\right)\bar{\psi}(u).$$

Proof. Since $\deg(u) \geq 1$, there exists a $\{u, v\} \in E$ such that

$$\bar{\psi}(u) - \bar{\psi}(v) \geq \frac{1}{\deg(u)} L\bar{\psi}(u) \tag{9}$$

for if no such $\{u, v\} \in E$ exists, then

$$L\bar{\psi}(u) = \deg(u)\frac{L\bar{\psi}(u)}{\deg(u)} < \sum_{\{u,w\} \in E} (\bar{\psi}(u) - \bar{\psi}(w)) = L\bar{\psi}(u),$$

a contradiction. So, since $L\bar{\psi}(u) = \lambda\bar{\psi}(u)$, we obtain from (9)

$$\bar{\psi}(v) \leq \left(1 - \frac{\lambda}{\deg(u)}\right)\bar{\psi}(u).$$

□

Denote by Δ the maximum degree of a vertex in X. Since

$$\left(1 - \frac{\lambda}{\deg(u)}\right) \leq \left(1 - \frac{\lambda}{\Delta}\right)$$

for all nonisolated vertices $u \in V$, we obtain the upper bound

$$\bar{\psi}(v) \leq \left(1 - \frac{\lambda}{\Delta}\right)^m \bar{\psi}(u)$$

on the cost of the current configuration $v \in V$ after m steps of greedy local search that started from a nonisolated vertex $u \in V$. In other words, the cost of the current configuration must decay exponentially fast (in m) towards an average cost configuration.

If δ is the minimum difference in cost between neighbouring configurations, we obtain that an average cost configuration is reached in m steps, where

$$m < \frac{\log\delta - \log\bar{\psi}(u)}{\log\left(1 - \frac{\lambda}{\Delta}\right)}.$$

6. The discrete nodal domain theorems

Let $X = (V, E)$ be a graph, and suppose $y \in \mathbb{R}^{|V|}$. A *strong positive (negative) sign graph* of y on X is a maximal, connected subgraph of X whose vertices u satisfy $y_u > 0$ ($y_u < 0$). A *weak positive (negative)* sign graph of y on X is a maximal, connected subgraph of X whose vertices u satisfy $y_u \geq 0$ ($y_u \leq 0$).

Let L be a real symmetric matrix with nonpositive off-diagonal elements. Suppose the rows and columns of L are indexed by I. Associate with L a graph Γ_L with vertex set I and edge set defined by $\{u, v\} \in E$ if and only if $A[u, v] < 0$. (In particular, we may take L to be the Laplacian of a graph X, in which case $\Gamma_L = X$.)

The discrete nodal domain theorems (which appear as one theorem below) bound the number of weak and strong sign graphs of eigenvectors of L on Γ_L:

Theorem 6.1 ([2]). *Let $\lambda_1 \leq \lambda_2 \cdots \leq \lambda_{|I|}$ be the eigenvalues of L, and suppose λ_m satisfies*

$$\lambda_{m-1} < \lambda_m = \lambda_{m+1} \cdots = \lambda_{m+r-1} < \lambda_{m+r}.$$

Then, any eigenvector associated with λ_m has at most $m + r - 1$ strong sign graphs and at most m weak sign graphs on Γ_L.

In particular, if we let $L = L(Q_2^n)$, then we obtain the following corollary:

Corollary 6.2. *The zero-mean cost function $\bar{\psi}$ of any NAESAT instance over n variables has at most $1 + n + \binom{n}{2}$ strong sign graphs, and at most $n + 2$ weak sign graphs on Q_2^n.*

Unfortunately, the previous corollary does not imply anything about the local minima and maxima of $\bar{\psi}$ on Q_2^n.

References

[1] L. Babai. Spectra of Cayley graphs. *J. Combin. Theory Ser. B* **27** (1979) 180–189.

[2] E. B. Davies, G. M. L. Gladwell, J. Leydold, and P. F. Stadler. Discrete nodal domain theorems. *Linear Algebra Appl.* **336** (2001) 51–60.

[3] R. Diestel. *Graph Theory, Second Edition.* Graduate Texts in Mathematics 173, Springer-Verlag, New York, 2000.

[4] M. R. Garey and D. S. Johnson. *Computers and Intractability: A Guide to The Theory of NP-Completeness.* A Series of Books in the Mathematical Sciences, W. H. Freeman and Co., San Francisco, 1979.

[5] C. Godsil and G. Royle. *Algebraic Graph Theory.* Graduate Texts in Mathematics 207, Springer-Verlag, New York, 2001.

[6] W. H. Greub. *Linear Algebra, Third Edition.* Die Grundlehren der mathematischen Wissenschaften, Band 97. Springer-Verlag, Berlin, 1967.

[7] L. K. Grover. Local search and the local structure of NP-complete problems. *Oper. Res. Lett.* **12** (1992) 235–243.

[8] G. James and A. Kerber. *The Representation Theory of the Symmetric Group.* Encyclopedia of Mathematics and Its Applications 16, Addison-Wesley, Reading, Massachusetts, 1981.

[9] C. M. Reidys and P. F. Stadler. Combinatorial landscapes. *SIAM Review*, to appear.

[10] D. Rockmore, P. Kostelec, W. Hordijk, P. F. Stadler. Fast Fourier transform for fitness landscapes. *Appl. Comput. Harmon. Anal.* **12** (2002) 57–76.