Partial Neighborhoods of Elementary Landscapes

L. Darrell Whitley

Andrew M. Sutton

Department of Computer Science
Colorado State University
Fort Collins, CO 80523
{whitley,sutton}@cs.colostate.edu

ABSTRACT

This paper introduces a new component based model that makes it relatively simple to prove that certain types of landscapes are elementary. We use the model to reconstruct proofs for the Traveling Salesman Problem, Graph Coloring and Min-Cut Graph Partitioning. The same model is then used to efficiently compute the average values over particular partial neighborhoods for these same problems. For Graph Coloring and Min-Cut Graph Partitioning, this computation can be used to focus search on those moves that are most likely to yield an improving move, ignoring moves that cannot yield an improving move. Let x be a candidate solution with objective function value $f(x)$. The mean value of the objective function over the entire landscape is denoted $\bar{f}$. Normally in an elementary landscape one can only be sure that a neighborhood includes an improving move (assuming minimization) if $f(x) > \bar{f}$. However, by computing the expected value of an appropriate partial neighborhood it is sometimes possible to know that an improving move exists in the partial neighborhood even when $f(x) < \bar{f}$.

Categories and Subject Descriptors

I.2.8 [**Artificial Intelligence**]: Problem Solving, Control Methods, and Search

General Terms

Theory, Algorithms

Keywords

Fitness Landscapes, Elementary Landscapes

1. INTRODUCTION

The *fitness landscape* for a combinatorial problem instance is defined by a triple (X, N, f). The *objective function* f maps $f : X \mapsto \mathbb{R}$ and without loss of generality we can define f so as either to be minimized or maximized over X.

GECCO'09, July 8–12, 2009, Montréal Québec, Canada.

We define a *neighborhood operator* as a function N that maps candidate solutions in X to subsets of X (elements of the power set) $N : X \mapsto \mathcal{P}(x)$. Given a candidate solution $x \in X$, $N(x)$ is the set of points reachable from x in one application of the neighborhood operator.

Thus the triple (X, N, f) defines a set of candidate solutions X, the neighborhood operator $N(x)$ which imposes a connective structure on the candidate solution points, and the objective function f assigns a value to each point. The fitness landscape provides a natural framework for the analysis of local search.

Elementary landscapes are a special class of fitness landscapes. One way to define elementary landscapes are as fitness landscapes where the neighborhood operator can be characterized by a *wave equation*. For all elementary landscapes it is possible to compute $\bar{f}$, the average solution evaluation over the entire search space. The wave equation also makes it possible to compute $\text{avg}\{f(y)\}_{y\in N(x)}$, the average value of the fitness function f evaluated over all of the neighbors of x:

$$\underset{y\in N(x)}{\text{avg}}\{f(y)\} = \frac{1}{|N(x)|}\sum_{y\in N(x)} f(y)$$

Other properties also follow. Assuming $f(x) \neq \bar{f}$ then

$$f(x) < \underset{y\in N(x)}{\text{avg}}\{f(y)\} < \bar{f} \quad \text{or} \quad f(x) > \underset{y\in N(x)}{\text{avg}}\{f(y)\} > \bar{f}.$$

This means that all maxima are greater than $\bar{f}$ and all minima are less than $\bar{f}$ [2]. Finally, $f(x) = \bar{f} \iff f(x) = \text{avg}\{f(y)\}_{y\in N(x)}$.

Grover [5] originally made the observation that there exists neighborhoods for the Traveling Salesman Problem, Graph Coloring, Min-Cut Graph Partitioning, Weight Partition, as well as Not-all-equal-Sat (NAES) that can be modeled using the wave equation.

This wave equation holds when the objective function f is an eigenfunction of the Laplacian of the graph induced by the neighborhood operator. Stadler [7] named this class of problems "elementary landscapes." Barnes et al. [1] have extended the notion of elementary landscapes to non-symmetric and non-regular neighborhoods. It can also be shown that a landscape with a symmetric neighborhood operator is elementary if and only the time series generated by a random walk on the landscape is an AR(1) process [8] [4].

For all of the elementary landscapes we have examined, the "components" that make up a solution x can be decomposed. In the Traveling Salesman Problem and Min-Cut Graph Partitioning and Graph Coloring, the components are the weights of a lower triangular cost matrix. In Graph Coloring the cost is usually 1 for each conflicted edge. In each of these problems, a uniform sample over all components in the cost matrix is used in the computation of $\bar{f}$. This is because over all points in the search space, all components in the cost matrix are uniformly sampled.

For any incumbent search point x, the cost matrix can be separated into those components that contribute to $f(x)$ and those components that do not contribute to $f(x)$. For example, for a tour x in the Traveling Salesman Problem the cost matrix can be broken down into those weights that contribute to $f(x)$ and those weights that do not contribute to $f(x)$. In Graph Coloring, the components are the set of edges and these can be decomposed into all those edges in x with cost 1 that contribute to $f(x)$ and the edges with cost 0 that do not contribute to $f(x)$. We can also characterize a neighbor $y \in N(x)$ in terms of the components that is shares with x and the components in y that are not found in x.

Let C represent the set of components that make up the cost function. If the components of C are uniformly sampled when every point in the search space is sampled, then $\bar{f}$ is always computable.

For the wave equation to hold, the components in x and the components in $N(x)$ must include all of the components in C. By a slight abuse of notation, we will let $(C-x)$ refer to components in C that do not appear in solution x or otherwise do not contribute to $f(x)$. When transforming x to some $y \in N(x)$ we subtract a subset of components from x and add new components from $(C-x)$ to the remaining components of x to create y. All components in x are uniformly sampled for potential removal from x. All components in $(C-x)$ are uniformly sampled in the set of neighbors $N(x)$. We show that any local search operator that obeys these principles induces an elementary landscape.

This breakdown of components into those that contribute to $f(x)$ and those that do not contribute to $f(x)$ makes it easy and intuitive to understand many of the fundamental properties that characterize elementary landscapes.

This paper also introduces *partial neighborhoods* of elementary landscapes. Local search methods do not necessarily evaluate all neighbors. For example, on the Traveling Salesman Problem some edges are so costly that they can never be part of the globally optimal solution. When one examines the neighborhood that induces an elementary landscape for the Graph Coloring problem or the Min-Cut Graph Partitioning problem, some neighborhood includes moves that can be predetermined to be non-improving moves: any reasonable search would eliminate these moves.

Let $N'(x)$ be a partial neighborhood such that $N'(x) \subset N(x)$. The landscape that is induced is generally not elementary. However, it is possible to define $N'(x)$ dynamically in such a way so that it is possible to compute $\text{avg}\{f(y)\}_{y \in N'(x)}$. This information can be valuable because $\text{avg}\{f(y)\}_{y \in N'(x)}$ can be computed much more efficiently than actually evaluating all the points in $N'(x)$.

To summarize, this paper presents a component based model that can be used to show that certain landscapes are elementary landscapes. The paper then focuses on the analysis of partial neighborhoods of elementary landscapes. We construct examples of partial neighborhoods denoted by $N'(x)$ that do not induce an elementary landscape, but where it is still possible to compute $\text{avg}\{f(y)\}_{y \in N'(x)}$.

2. ELEMENTARY LANDSCAPES

Let X be a set of solutions, $f : X \to \mathbb{R}$ be a fitness function, and $N : X \to \mathcal{P}(x)$ be a neighborhood operator. We can represent the neighborhood operator by its *adjacency matrix*

$$\mathbf{A}_{xy} = \begin{cases} 1 & \text{if } y \in N(x) \\ 0 & \text{otherwise} \end{cases}$$

The *degree matrix* $\mathbf{D}$ is defined as the diagonal matrix

$$\mathbf{D}_{xy} = \begin{cases} |N(x)| & \text{if } x = y \\ 0 & \text{otherwise} \end{cases}$$

Since a discrete function over the set of candidate solutions $g : X \to \mathbb{R}$ can be characterized as a vector in $\mathbb{R}^{|X|}$, any $|X| \times |X|$ matrix can be used as a linear operator on that function. The *Laplacian operator* is defined as

$$\Delta = \mathbf{A} - \mathbf{DI}$$

The Laplacian acts on the fitness function f as follows

$$\Delta f = \begin{bmatrix} \sum_{y \in N(x_1)} (f(y) - f(x_1)) \\ \sum_{y \in N(x_2)} (f(y) - f(x_2)) \\ \vdots \\ \sum_{y \in N(x_{|X|})} \left(f(y) - f(x_{|X|})\right) \end{bmatrix}$$

The element of this matrix-vector product corresponding to point x can thus be written as

$$\Delta f(x) = \sum_{y \in N(x)} (f(y) - f(x)) \tag{1}$$

In this paper, we will restrict our attention to *regular* neighborhoods, where $|N(x)| = d$ for a constant d for all $x \in X$. When a neighborhood is regular, $\Delta = \mathbf{A} - d\mathbf{I}$.

Stadler defines the class of *elementary landscapes* where the fitness function f is an eigenfunction of the Laplacian [8] (up to an additive constant). In particular, Grover's wave equation can be written as

$$\Delta f + k(f - \bar{f}) = 0$$

where $\bar{f}$ is the mean fitness value in X and k is a positive constant over the search space. It is easy to see that

$$\Delta f(x) = k\bar{f} - kf(x)$$

We can use this equation to express the *average fitness* across the neighborhood of any given candidate solution x. We denote this average fitness as

$$\begin{aligned}
\underset{y \in N(x)}{\text{avg}}\{f(y)\} &= \frac{1}{d} \sum_{y \in N(x)} f(y) \\
&= \frac{1}{d} \left(\sum_{y \in N(x)} f(y) - f(x) \right) + f(x) \\
&= \frac{1}{d} \Delta f(x) + f(x) && \text{by Eq. (1)} \\
&= f(x) + \frac{k}{d}(\bar{f} - f(x))
\end{aligned}$$

Thus, on an elementary landscape, the average fitness over the neighborhood of x can be completely characterized by an expression involving only the fitness of x, the average fitness $\bar{f}$, and a constant factor.

2.1 A Component Based Model

We will construct a component based model that can be used to characterize a neighborhood structure. To be more specific, given a point x, and its evaluation $f(x)$ and the mean fitness over all the points in a search space denoted by $\bar{f}$, we will compute $\text{avg}\{f(y)\}_{y\in N(x)}$.

The model looks at these calculations as a decomposition of the components of the evaluation function. Looking at landscapes in this manner makes it simpler to explore new questions about elementary landscapes. In this model, the neighborhood size is regular and denoted by d.

In order to have a component based model of an elementary landscape we need to define the set of components, denoted by C, that are used to construct the cost function. Intuitively, C is a collection of real or integer numbers; subsets of C are used to compute the value of $f(x)$ for a given point x. There are also 3 ratios p_1, p_2 and p_3 that are used in the following equations.

$$\bar{f} = p_3 \sum_{c\in C} c \quad \text{and therefore} \quad \sum_{c\in C} c = 1/p_3 \bar{f}$$

$$\begin{aligned}\underset{y\in N(x)}{\text{avg}}\{f(y)\} &= f(x) - p_1 f(x) + p_2((\sum_{c\in C} c) - f(x)) \\ &= f(x) - p_1 f(x) + p_2((1/p_3\bar{f}) - f(x))\end{aligned}$$

where $0 < p_1 < 1$ is the proportion of components that contribute to the evaluation of $f(x)$ and that change when a move is made; $0 < p_2 < 1$ is the proportion of components in $(C-x)$ that change when a move is made. Both p_1 and p_2 can be expressed relative to d, the size of the neighborhood. Finally, $0 < p_3 < 1$ is the proportion of the total components in C that contribute to the cost function for any randomly chosen solution; p_3 is independent of the neighborhood size.

Theorem 1. *If p_1, p_2 and p_3 can be defined for any regular landscape such that the evaluation function can be decomposed into components where $p_1 = \alpha/d$ and $p_2 = \beta/d$ and*

$$\bar{f} = p_3 \sum_{c\in C} c = \frac{\beta}{\alpha+\beta} \sum_{c\in C} c$$

then the landscape is elementary.

Proof.

$$\begin{aligned}\underset{y\in N(x)}{\text{avg}}\{f(y)\} &= f(x) - p_1 f(x) + p_2((\sum_{c\in C} c) - f(x)) \\ &= f(x) - p_1 f(x) + p_2((1/p_3\bar{f}) - f(x)) \\ &= f(x) - (p_1+p_2) f(x) + (p_2/p_3)\bar{f} \\ &= f(x) - \frac{\alpha+\beta}{d} f(x) + \frac{\beta/d}{\beta/(\alpha+\beta)}\bar{f} \\ &= f(x) + \frac{\alpha+\beta}{d}(\bar{f} - f(x)) \\ &= f(x) + \frac{k}{d}(\bar{f} - f(x))\end{aligned}$$

□

Note that p_1, p_2 and p_3 must be constants and

$$p_1 + p_2 = p_2/p_3 = k/d$$

where d is the size of the neighborhood and k is a constant. Thus when any two proportions are known, the third is automatically determined.

Because it is convenient to reason about all of the components that make up the cost function, as well as those components that play a role in the evaluation of $f(x)$ and it neighbors, it is useful to note that

$$\begin{aligned}d \cdot \underset{y\in N(x)}{\text{avg}}\{f(y)\} &= d\cdot f(x) + k(\bar{f} - f(x)) \\ &= (d-k)f(x) + kp_3(\sum_{c\in C} c) \\ &= (d-k)f(x) + \beta(\sum_{c\in C} c) \\ &= (d-k+\beta)f(x) + \beta(\sum_{c\in C} c - f(x))\end{aligned}$$

This computation also can be expressed as a 2-dimensional matrix M with d rows and $|C|$ columns. Each row represents a neighbor. Each column represents a component. The $m_{i,j}$ element of M stores the cost of the j^{th} component if that component appears in neighbor i; otherwise $m_{i,j} = 0$.

This view takes into account what happens to all components that contribute to the evaluation function as opposed to what happens on average or in expectation. We can also select particular rows of this matrix to create a partial neighborhood.

2.2 The Traveling Salesman Problem

We start with the Traveling Salesman Problem (TSP) as an example. The neighborhood is generated using 2-opt. Grover's original paper proved that an "exchange operator" induces an elementary landscape. Under the exchange operator city i exchanges positions with city j over all pairs of cities. Stadler showed that an "inversion operator" which is a superset of the 2-opt operator induces an elementary landscape.

Let E denote the set of all edges between cities (vertices) in the graph. Let $w_{i,j}$ be the weight (or distance) associated with edge $e_{i,j}$. This set of weights make up the set of components C, where $|C| = |E| = n(n-1)/2$. Note that this counts the number of weights in a lower triangular cost matrix for the TSP.

We first compute $\bar{f}$ and p_3. Since there are n edges in a given solution, it follows that

$$p_3 = \frac{n}{|C|} = \frac{n}{n(n-1)/2} = \frac{2}{n-1}$$

$$\bar{f} = p_3 \sum_{c\in C} = \frac{2}{n-1} \sum_{e_{i,j}\in E} w_{i,j}$$

where $\sum_{e_{i,j}\in E} w_{i,j}$ counts each edge only once.

To compute p_1 note there are n edges in any solution, and 2-opt changes exactly 2 edges. Therefore

$$p_1 = 2/n \quad \text{and} \quad p_1 = \frac{2(n-3)/2}{n(n-3)/2} = \frac{\alpha}{d}$$

To compute p_2 note there are $|C| - n$ edges in C with the edges in $f(x)$ removed, and 2 new edges are picked from these edges. Therefore

$$p_2 = \frac{2}{n(n-1)/2 - n} = \frac{2}{n(n-3)/2} = \frac{\beta}{d}$$

Adding the terms to the component model yields:

$$\begin{aligned}\operatorname*{avg}_{y\in N(x)}\{f(y)\} &= f(x) - p_1 f(x) + p_2(1/p_3\bar{f} - f(x)) \\ &= f(x) - \frac{2}{n}f(x) + \frac{2}{n(n-3)/2}[(\frac{n-1}{2}\bar{f}) - f(x)] \\ &= f(x) + \frac{n-1}{n(n-3)/2}(\bar{f} - f(x))\end{aligned}$$

where $k = n-1$ and the neighborhood size is $d = n(n-3)/2$.

It should be noted that most of the work on elementary landscapes has focused on the Traveling Salesman Problem (TSP). This includes a proof that the symmetric Traveling Salesman Problem is elementary under 2-exchange [5], and 2-opt and 3-exchange, [2]. The antisymmetric Traveling Salesman Problem under 2-opt and 2-exchange [8] is also elementary, as is the weakly-symmetric Traveling Salesman Problem [6], and variants of the multiple Traveling Salesman Problem [3].

2.3 Min-Cut Graph Partitioning

Let G denote a graph with weighted edges and n vertices (where n is even). The Min-Cut Graph Partitioning problem is to find a partition of the vertices in G into a left hand side (LHS) and a right hand side (RHS) such that $|\text{LHS}| = |\text{RHS}| = n/2$ and the weighted sum of the edges connecting vertices in the LHS and RHS partitions is minimized.

The neighborhood operator will exchange every vertex in the LHS with every vertex in the RHS. There are $n/2$ vertices in the LHS and $n/2$ in the right hand side. Thus there are $(n/2)^2 = n^2/4$ neighbors.

Let E denote the set of all edges in graph G. Let $w_{i,j}$ be the weight (or distance) associated with edge $e_{i,j} \in E$. This set of weights make up the set of components C, where $|C| = |E| = n(n-1)/2$.

Consider an arbitrary edge $e_{i,j}$. Grover notes that if vertex i is one side, then the fraction of total configurations where j is on the other side is $\frac{n/2}{n-1}$. To state this more clearly, each vertex v is connected to $n-1$ other vertices; if $v \in$ LHS then there are $(n/2)$ edges from v to edges in RHS. We impose an order on the vertices and weights and only count them once, so that we obtain:

$$p_3 = \frac{n/2}{n-1} = \frac{n^2/4}{|C|}$$

$$\bar{f} = p_3 \sum_{c\in C} c = \frac{n}{2(n-1)} \sum_{e_{i,j}\in E} w_{i,j}$$

To compute p_1 let edge $e_{l,r}$ denote an edge with vertices $l \in$ LHS and $r \in$ RHS which contributes to $f(x)$. We can move l to the RHS and keep r in the RHS and move any of the other $(n/2 - 1)$ vertices in the RHS to the LHS. By symmetry we can move r to the LHS and keep l in the LHS in $(n/2 - 1)$ ways. Combining these moves, there are $2(n/2 - 1)$ ways to remove the contribution of weight $w_{l,r}$ from the current solution $f(x)$. Since there are $n^2/4$ total neighbors:

$$p_1 = \frac{2(n/2-1)}{n^2/4} = \frac{n-2}{n^2/4} = \frac{\alpha}{d}$$

Let $e_{a,b}$ be an edge from C that does not contribute to $f(x)$. This means vertices a and b must be on the same side of the partition. The weight associated with this edge adds to the cost function if vertex a changes sides or if vertex b changes sides. There are $n/2$ vertices on the opposite side with which each can exchange positions. Thus there are $n/2 + n/2 = n$ ways that an arbitrary edge that does not contribute to $f(x)$ can contribute its weight to the cost of a neighbor of x. Since there are $n^2/4$ total neighbors:

$$p_2 = \frac{n}{n^2/4} = \frac{\beta}{d}$$

Adding these terms to the component model yields:

$$\begin{aligned}\operatorname*{avg}_{y\in N(x)}\{f(y)\} &= f(x) - p_1 f(x) + p_2(1/p_3(\bar{f} - f(x))) \\ &= f(x) - \frac{n-2}{n^2/4}f(x) + \frac{n}{n^2/4}\left[\frac{2(n-1)}{n}\bar{f} - f(x)\right] \\ &= f(x) + \frac{2(n-1)}{n^2/4}(\bar{f} - f(x))\end{aligned}$$

where $k = 2(n-1)$ and the neighborhood size is $d = n^2/4$. Grover simplifies this to obtain:

$$\operatorname*{avg}_{y\in N(x)}\{f(y)\} = f(x) + \frac{8(n-1)}{n^2}(\bar{f} - f(x))$$

This construction is significantly simpler than Grover's original proof (which also contains a notation error in equation 1.1.3).

2.4 Graph Coloring

Let G be a graph, V the set of vertices, and E the set of edges. The graph coloring problem involves assigning one of r number of colors to the vertices of a graph. Given r colors, the goal is to find a coloring that minimizes the number of conflicts. A conflict exists if two adjacent vertices have the same color. The evaluation function $f(x)$ simply counts how many adjacent vertices have the same color. This means that every edge either 1) contributes 1 to the cost function if the vertices connected by that edge have the same color, or 2) contributes 0 to the cost function if the vertices connected by the edge have a different color. Here the set of components C corresponds directly to the set of edges; this time, each component has a "weight" contribution of exactly 1.

Edges that have a cost of 1 contribute to $f(x)$. The cost function $f(x)$ counts the number of edges in solution x that have cost 1. Thus, for solution x there are $|E| - f(x)$ edges that have a cost of 0; these are the edges that are in C and do not contribute to $f(x)$. Note that $\sum_{c\in C} c = |E|$.

The neighborhood operator is to recolor a vertex in the graph. Since there are $|V|$ vertices, and each vertex can be recolored in $r-1$ ways, the size of the neighborhood is $|V|(r-1)$. The average cost over all solutions will be

$$\bar{f} = 1/r \sum_{c\in C} c = |E|/r$$

This is because for every edge, once one vertex is colored, the probability the second vertex associated with an edge is colored so as to yield conflict is $1/r = p_3$.

Consider two vertices v_1 and v_2 that are the same color and connected by an edge. There are $r-1$ colors that can be assigned to either v_1 or v_2 that will remove the conflict. Thus, there are $2(r-1)$ assignments removing each conflict and it follows that:

$$p_1 = 2(r-1)/(|V|(r-1)) = \frac{\alpha}{d}$$

When a conflict does not exist, there are exactly two ways for the conflict to be generated. Given a vertex v_1 and v_2 that are different colors and connected by an edge, either v_1 is colored the same as v_2, or v_2 is colored the same as v_1. There are only 2 ways this can happen and it follows that:

$$p_2 = 2/(|V|(r-1)) = \frac{\beta}{d}$$

Adding the terms to the component model yields:

$$\begin{aligned}\operatorname*{avg}_{y\in N(x)}\{f(y)\} &= f(x) - p_1 f(x) + p_2(1/p_3\bar{f} - f(x)) \\ &= f(x) - \tfrac{2(r-1)}{|V|(r-1)} f(x) + \tfrac{2}{|V|(r-1)}(r\bar{f} - f(x)) \\ &= f(x) + \frac{2r}{|V|(r-1)}(\bar{f} - f(x))\end{aligned}$$

This final form satisfies Grover's wave equation where $k = 2r$ and neighborhood size $d = |V|(r-1)$.

3. PARTIAL NEIGHBORHOODS WITHIN ELEMENTARY LANDSCAPES

Can an elementary landscape exist within a larger elementary landscape? In particular, can a partial neighborhood of an elementary landscape also be an elementary landscape? The answer seems to be yes, but not usually. We have found fascinating examples for small fixed size problems where an elementary landscape could be broken into subspaces, each of which is elementary. However, we have not found this to be true for arbitrary size problems.

But we can also ask another question. Given a partial neighborhood $N'(x)$ specific to point x, can we calculate $\text{avg}\{f(y)\}_{y\in N'(x)}$? When we restrict moves (and therefore neighbors), the components that make up the neighbors of x changes. One advantage to our component based model is that it can be applied locally even when the landscape is not elementary.

3.1 A Partial Neighborhood for the TSP

As an illustration of partial neighborhoods, we define the "Diagonal-2-opt" neighborhood for the Traveling Salesman Problem. This neighborhood requires that the number of cities be even. It should be noted that this neighborhood is (perhaps) the most artificial of the partial neighborhoods that are presented in this paper. But it has some special properties that are informative.

Let $\mathcal{D}(x)$ be the set of diagonal edges relative to x. Assuming n (the number of cities) is even, $|\mathcal{D}(x)| = n/2$. For a given solution x these edges are defined as follows:

$$e_{v_i,(v_{i+(n/2)})} \in \mathcal{D}(x) \quad \forall i = 1...n/2$$

We also compute the sum $D(x)$ over the corresponding set of weights.

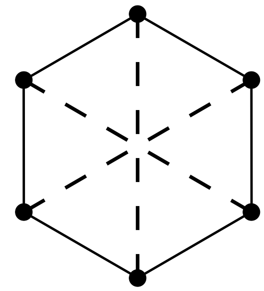

Figure 1: A graph on six vertices with tour edges (solid) and interior diagonal edges (dashed).

$$D(x) = \sum_{i=1}^{n/2} w_{v_i,(v_{i+(n/2)})}$$

A graph with six vertices along with the tour edges and diagonal edges is illustrated in figure 1. We will define a partial neighborhood around x such that only diagonal-2-opt moves are allowed that cut adjacent edges and cause diagonal edges to become adjacent edges in the neighboring solutions: thus, all of the solutions are made up of edges in solution x unioned with $\mathcal{D}(x)$.

The sum of all the edges that make up the diagonal neighborhood with respect to x is given by $f(x) + D(x)$. This defines a local version of C' particular to x. We can also *potentially* define a local average, denoted by $\bar{f}^x$ that depends on x and the corresponding partial diagonal-2-opt neighborhood. In some partial neighborhoods, $\bar{f}^x$ and p_3 are not well defined. However, in this case we can show that $\bar{f}^x$ and p_3 are well defined for a specific TSP instance. Therefore:

$$\bar{f}^x = \frac{n}{n+(n/2)}(f(x) + D(x)) = \frac{f(x) + D(x)}{1.5}$$

where $p_3 = \frac{n}{|C'|} = \frac{n}{n+(n/2)} = 2/3$. This implies that $D(x) = 1.5\bar{f}^x - f(x)$. Let $\text{avg}\{f(y)\}_{y\in N'(x)}$ denote the average value of the diagonal neighbors of x. In this case, $p_1 = 2/n$ does not change, but p_2 now selects 2 edges from $\mathcal{D}$ which has $n/2$ elements. Thus $p_2 = \frac{2}{n/2} = 4/n$. Adding the terms to the component model yields:

$$\begin{aligned}\operatorname*{avg}_{y\in N'(x)}\{f(y)\} &= f(x) - p_1 f(x) + p_2(1/p_3\bar{f}^x - f(x)) \\ &= f(x) - \frac{2}{n} f(x) + \frac{2}{n/2}(1.5\bar{f}^x - f(x)) \\ &= f(x) + \frac{3}{n/2}(\bar{f}^x - f(x))\end{aligned}$$

The final expression is similar in form to Grover's wave equation with $k = 3$ and neighborhood size d' $= n/2$. However, does this define an elementary landscape?

It turns out that a TSP with 6 cites is a special case. Each diagonal neighborhood is closed under diagonal-2-opt and defines a subspace of 6 tours. The entire search space contains $5!/2 = 60$ tours. The entire search space is an elementary landscape under 2-opt, but it also breaks into 10 distinct subspaces each of which is an elementary landscape under diagonal-2-opt. Each subspace contains 6 tours, where

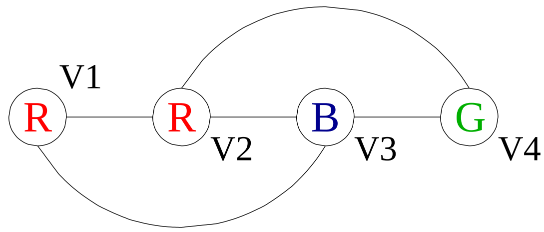

Figure 2: A Graph Coloring problem with vertices $V = \{v_1, v_2, v_3, v_4\}$.

these 6 tours define a group closed under the diagonal-2-opt neighbors. Since the search space contains only 60 tours, these observations can be proven by enumeration.

When $n > 6$, the above equation still holds for computing the average of the partial neighborhood, but $\bar{f}^x$ now is a local average that depends on x. Therefore, the resulting subspace and landscape is not elementary. Nevertheless, we can still locally compute $\text{avg}\{f(y)\}_{y \in N'(x)}$ for the partial neighborhood defined by the diagonal-2-opt neighborhood.

We have also examined the diagonal-2-opt neighborhood for instances where n > 6. In general, there appears to be a group closure over a subset of permutations that are reachable by diagonal-2-opt moves. But in general, these are not elementary landscapes.

It should also be pointed out that a local average can be computed over other partial neighborhoods of the Traveling Salesman Problem. These neighborhoods must be selected so that p_2 can be meaningfully defined with respect to C', but the other information needed to construct the local average are trivial to derive.

3.2 Partial Neighborhoods for Graph Coloring

For Graph Coloring we explore a slightly different question with regard to partial neighborhoods.

When searching the Graph Coloring neighborhood, certain moves do not appear to be reasonable if we are interested in removing conflicts. (As we will show, the same is true for Min-Cut Graph Partitioning.)

A Graph Coloring problem is given in figure 2 where

$$V = \{v_1, v_2, v_3, v_4\} \text{ and } E = \{e_{1,2},\ e_{2,3},\ e_{3,4},\ e_{1,3},\ e_{2,4}\}.$$

In figure 2 vertices v_1 and v_2 are the same color R, v_3 is color B, and v_4 is color G. The only edge in conflict is $e_{1,2}$. Changing the color of v_1 and v_2 can remove the conflict, so recoloring v_1 and v_2 are reasonable neighborhood moves. But changing the color of v_3 and v_4 is useless: recoloring these vertices cannot remove existing conflicts, but will generate new conflicts.

Let $Degree(v)$ be a vector that stores the degree of every vertex $v \in |V|$. Let Q_x be the set of vertices such that if edge $e_{i,j}$ contributes cost to $f(x)$ then vertices i and j are members of set Q_x.

Theorem 2. *For the partial neighborhood denoted by $N'(x)$ such that only vertices in Q_x are recolored, the partial neighbor for the Graph Coloring problem is given by:*

$$\underset{y \in N'(x)}{\text{avg}}\{f(y)\} = f(x) + \frac{[(\sum_{v \in Q_x} Degree(v)) - (2r)f(x)]}{|Q_x|(r-1)}$$

Proof. We first consider the full component based model for Graph Coloring.

$$\underset{y \in N(x)}{\text{avg}}\{f(y)\} = f(x) - \frac{2(r-1)}{|V|(r-1)} f(x) + \frac{2}{|V|(r-1)}(|E| - f(x))$$

Only vertices in Q_x will be recolored, thus, $|V|$ changes to $|Q_x|$ in the equations. The number of ways a vertex can remove conflicts (α) does not change, therefore:

$$p_1 = \frac{2(r-1)}{|Q_x|(r-1)}$$

We are interested in the subset of the edges in E that change in the partial neighborhood. That is, we can re-express

$$r\bar{f} = |E| = |E_x + E_z|$$

where E_x are the edges we wish to keep in our partial neighborhood and E_z are the edges we wish to delete.

For all edges $e_{i,j}$ that contributes cost to $f(x)$ the vertices i and j are members of set Q_x. E_x therefore includes all edges that are incident on the vertices in Q_x. We can determine which elements are in E_x as a by-product of evaluating $f(x)$. Recall $Degree(v)$ is a vector that stores the degree of every vertex v in graph G. It follows that:

$$|E_x| = \sum_{v \in Q_x} Degree(v) - f(x)$$

The $-f(x)$ term is due to the fact that by counting the degree of each vertex the sum includes the edges that contribute to $f(x)$ twice. The remainder of the edges that touch a vertex involved in a conflict are members of E_x that are counted only once.

By construction, $(|E_x| - f(x))$ counts the number of edges such that $e_{i,j}$ where $i \in Q_x$ and $j \notin Q_x$ or $j \in Q_x$ and $i \notin Q_x$. Note this counts edges in the partial neighborhood that do not contribute to $f(x)$. In the full neighborhood, these edges have 2 vertices that are recolored, and $\beta = 2$. In this partial neighborhood, these edges have only 1 vertex that is recolored; thus there is only 1 way to generate a conflict, and $\beta = 1$.

$$p_2 = \frac{1}{|Q_x|(r-1)}$$

We now apply the component model to the partial neighborhood.

$$\begin{aligned}
\underset{y \in N'(x)}{\text{avg}}\{f(y)\} &= f(x) - \tfrac{2(r-1)}{|Q_x|(r-1)} f(x) + \tfrac{1}{|Q_x|(r-1)}(|E_x| - f(x)) \\
&= f(x) + \tfrac{|E_x| - f(x) - 2(r-1)f(x)}{|Q_x|(r-1)} \\
&= f(x) + \tfrac{(\sum_{v \in Q_x} Degree(v)) - 2f(x) - 2(r-1)f(x)}{|Q_x|(r-1)} \\
&= f(x) + \tfrac{[(\sum_{v \in Q_x} Degree(v)) - (2r)f(x)]}{|Q_x|(r-1)}
\end{aligned}$$

□

This partial neighborhood is not elementary because it is locally defined with respect to x. Nevertheless the average value of all of the neighbors in the dynamically defined neighborhood can be cheaply computed exploiting a

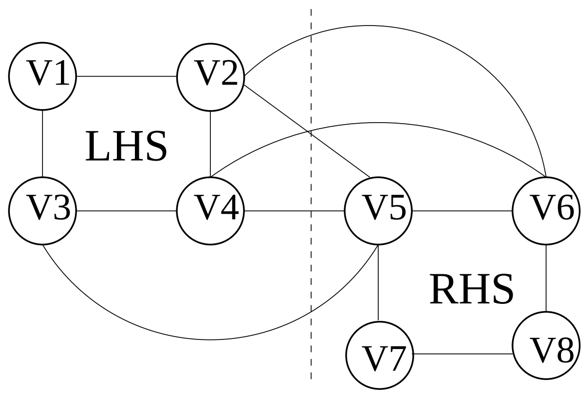

Figure 3: A Min-Cut Graph Partitioning problem with $LHS = \{v_1, v_2, v_3, v_4\}$ and $RHS = \{v_5, v_6, v_7, v_8\}$.

decomposition of the full Graph Coloring elementary landscape neighborhood. Since this partial neighbor removes all conflicts, but does not consider moves that cannot remove conflicts:

$$\operatorname*{avg}_{y \in N'(x)}\{f(y)\} \leq \operatorname*{avg}_{y \in N(x)}\{f(y)\}$$

Also note that

$$\sum_{v \in Q_x} Degree(v) < r + 2f(x) \quad \Longleftrightarrow \quad \operatorname*{avg}_{y \in N'(x)}\{f(y)\} < f(x)$$

and thus under these conditions an improving move is guaranteed to exist.

Empirically, we have observed cases where

$$\operatorname*{avg}_{y \in N'(x)}\{f(y)\} < f(x) < \operatorname*{avg}_{y \in N(x)}\{f(y)\} < \bar{f}$$

We have also observed cases where $\text{avg}\{f(y)\}_{y \in N'(x)} < f(x)$ when the global optimum is contained in the partial neighborhood $N'(x)$. Thus, calculations about partial neighborhoods can point to improving moves even when $f(x)$ is already near optimal (but not locally optimal) solution.

Also note that as $f(x)$ gets smaller, there are fewer conflicts and the it becomes cheaper to compute

$$\sum_{v \in Q_x} Degree(v).$$

3.3 Partial Neighborhoods for Min-Cut Graph Partitioning

Figure 3 shows a Min-Cut Graph Partitioning problem with 8 vertices and a partition where LHS = $\{v_1, v_2, v_3, v_4\}$ and RHS = $\{v_5, v_6, v_7, v_8\}$.

There are moves that are included in the neighborhood for the Min-Cut Graph Partitioning problem that also are ineffective. Without loss of generality, assume vertex v_1 is the LHS partition and has no edges that connect to the RHS partition; furthermore, assume vertex v_2 is in the RHS and has no edges that connect to the LHS partition. Then when swapping vertices in the LHS and RHS, it is unreasonable to swap the pair v_1 and v_2 because this cannot reduce the cost function relative to the current solution x, but instead must always increase the cost function.

Thus, the calculation of a partial neighborhood for Min-Cut Graph Partition can be constructed where

$$\operatorname*{avg}_{y \in N'(x)}\{f(y)\} \leq \operatorname*{avg}_{y \in N(x)}\{f(y)\}$$

Let n_L count the number of vertices in the LHS which have no edges that connect to the right hand size. Let n_R count the number of vertices in the RHS which have no edges that connect to the right hand size.

We again first start with the full neighbor component model for Min-Cut Graph Partitioning.

$$\operatorname*{avg}_{y \in N(x)}\{f(y)\} = f(x) - p_1 f(x) + p_2(1/p_3 \bar{f} - f(x))$$

$$\text{Note:} \quad \frac{1}{p_3}\bar{f} = \sum_{e_{i,j} \in E} w_{i,j}$$

$$\operatorname*{avg}_{y \in N(x)}\{f(y)\} = f(x) - \frac{n-2}{n^2/4} f(x) + \frac{n}{n^2/4}[(\sum_{e_{i,j} \in E} w_{i,j}) - f(x)]$$

The partial neighborhood removes swaps that cannot yield an improving move. Let d' denote the size of the new neighbor, and let W' represent the sum of all the weights that are eliminated when these moves are excluded.

THEOREM 3. *A partial neighborhood $N'(x)$ exists for the Min-Cut Graph Partitioning problems such that*

$$\operatorname*{avg}_{y \in N'(x)}\{f(y)\} = f(x) - \frac{2n-2}{d'} f(x) + \frac{n(\sum_{e_{i,j} \in E} w_{i,j}) - W'}{d'}$$

$$\textit{where } d' = n^2/4 - |n_L||n_r|$$

$$\textit{and } W' = |n_L| \sum_{i \in V, x \in n_R} w(i,x) + |n_R| \sum_{i \in V, x \in n_L} w(i,x)$$

PROOF. We start with the equations for the full neighborhood.

$$\begin{aligned}\operatorname*{avg}_{y \in N(x)}\{f(y)\} &= f(x) - \frac{n-2}{n^2/4} f(x) + \frac{n}{n^2/4}\left[(\sum_{e_{i,j} \in E} w_{i,j}) - f(x)\right] \\ &= f(x) - \frac{2n-2}{n^2/4} f(x) + \frac{n}{n^2/4} \sum_{e_{i,j} \in E} w_{i,j}\end{aligned}$$

The edges that are accounted for and removed by the term

$$((n-2)/d')f(x)$$

are all edges that cross from the LHS to the RHS and contribute to the cost function, and thus the $\alpha = (n-2)$ is not affected by vertices in n_L or n_R.

Each move involving edges that currently do not contribute to the cost function causes $(n/2) + (n/2) = n = \beta$ new edges to contribute to the cost function. This does not change under the partial neighborhood: if a move does not occur the associated weights simply do not contribute to the cost function.

We next view this as the aggregate over all the neighbors.

$$d \cdot \operatorname{avg}_{y \in N(x)}\{f(y)\} = (d-k)f(x) + \beta(\sum_{c \in C} c)$$
$$= (d-(2n-2))f(x) + n \sum_{e_{i,j} \in E} w_{i,j}$$

We next explicitly remove those edge weights that are removed from the neighborhood when the vertices in n_L and n_R do not swap positions. We also change the number of neighbors to d'. The equation still reduces to the wave equation for the full neighborhood for Min-Cut Graph Partitioning when $d' = d$ and $W' = 0$.

$$d' \cdot \operatorname{avg}_{y \in N'(x)}\{f(y)\} = (d' - (2n-2))f(x) + n \left(\sum_{e_{i,j} \in E} w_{i,j} \right) - W'$$

Note this is equivalent to:

$$\operatorname{avg}_{y \in N'(x)}\{f(y)\} = f(x) - \frac{2n-2}{d'} f(x) + \frac{n(\sum_{e_{i,j} \in E} w_{i,j}) - W'}{d'}$$

To complete the proof, we define d' and W'. Consider $l \in n_L$ and $r \in n_R$. A move that swaps the vertices (l, r) can never reduce the cost function, it can only increase the cost function. Thus all such moves can be removed from the neighborhood. There are $|n_L||n_R|$ moves that an be eliminated. Therefore:

$$d' = (n/2)^2 - |n_L||n_R|$$

W' is the sum of all the weights that are no longer contribute to the partial neighborhood. Every move that swaps vertices (l, r) with $l \in n_L$ and $r \in n_R$ is eliminated.

For each vertex in n_R this removes $|n_L|$ moves. For each vertex in n_L this removes $|n_R|$ moves. Thus the weights associated with edges incident on these vertices must be remove $|n_L|$ or $|n_R|$ times. Therefore

$$W' = |n_L| \sum_{i \in V, x \in n_R} w(i,x) + |n_R| \sum_{i \in V, x \in n_L} w(i,x)$$

□

Both d' and W' are conveniently calculated as a side effect of the evaluation of $f(x)$. A vector $Weights(x)$ can be precomputed that stores the sum of the weights associated with edges incident on vertex x.

The information needed to compute W' is found in the vector $Weights(i)$ since

$$\text{If } x \in n_L \text{ then } \quad \sum_{i \in V} w(i,x) = Weights(x)$$

$$\text{If } x \in n_R \text{ then } \quad \sum_{i \in V} w(i,x) = Weights(x)$$

Therefore

$$W' = |n_L| \sum_{x \in n_R} Weights(x) + |n_R| \sum_{x \in n_L} Weights(x)$$

Thus, using this method, the partial neighborhood average $\operatorname{avg}\{f(y)\}_{y \in N'(x)}$ is inexpensive to compute relative to the cost of exhaustively evaluating the entire partial neighborhood to calculate the average fitness. The partial neighborhood average also costs less to compute when there are fewer vertices involved in conflicts.

4. CONCLUSIONS

This paper has employed a component based model to reconstruct elementary landscape proofs for a number of combinatorial optimization problems. Furthermore, this paper has examined ways in which partial neighborhoods can be evaluated by exploiting knowledge about the elementary landscape structures of the Traveling Salesman Problem, Graph Coloring Problem, and Min-Cut Graph Partitioning Problem. These results are important because neighborhoods which restrict search to the most promising moves are more likely to be used by modern local search algorithms. Future work includes extending the component-based model to other landscapes.

5. ACKNOWLEDGMENTS

This research was sponsored by the Air Force Office of Scientific Research, Air Force Materiel Command, USAF, under grant number FA9550-08-1-0422. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright notation thereon.

6. REFERENCES

[1] J. W. Barnes, B. Dimova, and S. P. Dokov. The theory of elementary landscapes. *Applied Mathematics Letters*, 16(3):337–343, April 2003.

[2] B. Codenotti and L. Margara. Local properties of some NP-complete problems. Technical Report TR 92-021, International Computer Science Institute, Berkeley, CA, 1992.

[3] B. Colletti and J. W. Barnes. Linearity in the traveling salesman problem. *Applied Mathematics Letters*, 13(3):27–32, April 2000.

[4] B. Dimova, J. Wesley Barnes, and E. Popova. Arbitrary elementary landscapes & AR(1) processes. *Appl. Math. Lett.*, 18(3):287–292, 2005.

[5] L. K. Grover. Local search and the local structure of NP-complete problems. *Operations Research Letters*, 12:235–243, 1992.

[6] A. Solomon, J. W. Barnes, S. P. Dokov, and R. Acevedo. Weakly symmetric graphs, elementary landscapes, and the TSP. *Appl. Math. Lett.*, 16(3):401–407, 2003.

[7] P. F. Stadler. Toward a theory of landscapes. In R. Lopéz-Peña, R. Capovilla, R. García-Pelayo, H. Waelbroeck, and F. Zertruche, editors, *Complex Systems and Binary Networks*, pages 77–163. Springer Verlag, 1995.

[8] P F. Stadler. Landscapes and their correlation functions. *Journal of Mathematical Chemistry*, 20:1–45, 1996.