

Fast Fourier Transform for Fitness Landscapes

DAN ROCKMORE^a, PETER KOSTELEČ^a, WIM HORDIJK^b, PETER F. STADLER^{b,c,*}

^aDepartment of Computer Science, Dartmouth College
6211 Sudikoff Laboratory, Hanover, NH 03755-3510, USA
Phone: (603) 646 2206 Fax: (603) 646-1672
E-Mail: {rockmore,geelong}@cs.dartmouth.edu

^bThe Santa Fe Institute
1399 Hyde Park Road, Santa Fe, NM 87501, USA
Phone: (505) 984 8800 Fax: (505) 982 0565
E-Mail: {stadler,wim}@santafe.edu

^cInstitut für Theoretische Chemie, Universität Wien
Währingerstraße 17, A-1090 Wien, Austria
Phone: **43-1-4277-52737 Fax: **43-1-4277-52793
E-Mail: studla@tbi.univie.ac.at

*Address for correspondence

Abstract

We cast some classes of fitness landscapes as problems in spectral analysis on various Cayley graphs. In particular, landscapes derived from RNA folding are realized on Hamming graphs and analyzed in terms of Walsh transforms; assignment problems are interpreted as functions on the symmetric group and analyzed in terms of the representation theory of S_n . We show that explicit computations of the Walsh/Fourier transforms are feasible for landscapes with up to 10^8 configurations using Fast Fourier Transform techniques.

We find that the cost function of a linear sum assignment problem involves only the defining representation of the symmetric group, while quadratic assignment problems are superpositions of the representations indexed by the partitions (n) , $(n-1, 1)$, $(n-2, 2)$, and $(n-2, 1, 1)$. These correspond to the four smallest eigenvalues of the Laplacian of the Cayley graph obtained from using transpositions as the generating set on S_n .

Keywords. Spectral Analysis, Fast Fourier Transform, Walsh Functions, Cayley Graphs, Fitness Landscapes, Assignment Problems, RNA Folding.

AMS Subject Classification. 20C40, 42C10, 05C90.

1. Introduction

From the mathematical point of view, a landscape consists of three ingredients: (i) a set V of “configurations” which we shall assume to be finite but very large, (ii) a cost or fitness function $f : V \rightarrow \mathbb{R}$ that evaluates the configurations, and (iii) some sort of additional geometrical, topological, or algebraic structure \mathcal{X} on V that allows us to define notions of closeness, similarity, or dissimilarity among the configurations. The structure \mathcal{X} , which turns the set V into the *configuration space* (V, \mathcal{X}) , is determined by the particular application, e.g. a heuristic search procedure for a combinatorial optimization problem, or by the mechanisms of mutation and recombination in biological evolution.

A very promising approach in landscape theory is the decomposition of the fitness function $f : V \rightarrow \mathbb{R}$ in terms of a basis (of the vector space \mathbb{R}^V) that is induced in some natural way by \mathcal{X} . In other words, we search for a suitable spectral theory of the combinatorial space (V, \mathcal{X}) , which we then use to “Fourier transform” f with respect to a suitable set eigenfunctions of (V, \mathcal{X}) . The resulting “Fourier coefficients”, may reveal the important features of the landscape much more readily than f itself.

The simplest case is based on the notion of a *move set*. For each $x \in V$ we define the set $\mathcal{N}(x)$ of *neighbors* of x . The elements of $\mathcal{N}(x)$ are those configurations that can be reached from x in a single step. Many cases naturally suggest symmetric neighborhood relations, in which each step in the move set is *reversible*, defined as $x \in \mathcal{N}(y) \iff y \in \mathcal{N}(x)$. Thus, the set

$$\mathcal{E} = \{(x, y) \mid x \in V, y \in \mathcal{N}(x)\} \quad (1)$$

becomes the edge set of an undirected graph with vertex set V . We assume that this graph is loop-free, so that $x \notin \mathcal{N}(x)$. In many cases, there is a high intrinsic regularity in the way in which the move set is constructed. This leads to graphs with a very high degree of regularity. In particular, all configuration spaces considered in this contribution will give rise to regular graphs.

The most immediate algebraic representation of a graph Γ is its adjacency matrix $\mathbf{A}(\Gamma)$ with entries $\mathbf{A}_{xy} = 1$ if x and y are adjacent and $\mathbf{A}_{xy} = 0$ otherwise. In many respects it is more convenient to work with the *Laplacian matrix* $\mathbf{\Delta} = \mathbf{D} - \mathbf{A}$, where \mathbf{D} is the diagonal matrix of vertex degrees. For surveys on graph Laplacians see e.g. [1, 2, 3, 4]. For D -regular graphs, $\mathbf{D} = D\mathbf{I}$, hence, \mathbf{A} and $\mathbf{\Delta}$ have the same eigenvectors and their eigenvalues are related by $\Lambda_k^{\mathbf{\Delta}} = D - \lambda_k^{\mathbf{A}}$.

The eigenvalues of the Laplacian for a graph Γ are non-negative, with $\Lambda_0^{\mathbf{\Delta}} = 0$ belonging to a constant eigenvector. The eigenvectors of graphs share some important properties with the eigenfunctions of Laplacian operators on Riemannian manifolds [5, 6]. In particular, the number of “nodal domains”, i.e. connected components of Γ on which the eigenfunction does not change sign, generally increases with $\Lambda_k^{\mathbf{\Delta}}$. In addition, the cost functions of a number of prominent combinatorial optimization problems, among them the traveling salesman problem (TSP), graph bi-partitioning, and certain spin glass models, are eigenfunctions of graphs associated with search heuristics for these problems [7, 8, 9, 10, 11].

It seems natural, therefore, to try to extract information about fitness landscapes and cost functions of combinatorial optimization problems by expanding them in a

kind of “Fourier series”

$$f(x) = \sum_{i=0}^{|V|-1} a_i \varphi_i(x) \quad (2)$$

in terms of an orthonormal basis $\{\varphi_i\}_{i=0}^{|V|-1}$ of the Laplacian matrix Δ .

A general idea on a landscape’s ruggedness can be obtained by condensing the information of Fourier coefficients a_k into the so-called amplitude spectrum [10], defined as

$$\beta_p = \sum_{k: \Delta \varphi_k = \lambda_p \varphi_k} |a_k|^2. \quad (3)$$

or

$$B_p = \beta_p / \sum_{q \neq 0} \beta_q \quad (4)$$

in the normalized form. Note the $\sum_{q \neq 0} \beta_q = \sum_{x \in V} (f(x) - \bar{f})^2$ is the variance of the landscape. The amplitude spectrum thus measures the importance of the individual eigenspaces.

Provided the configuration space is “sufficiently regular”, explicit computations become feasible for interestingly large examples if *Fast Fourier Transform* (FFT) techniques are available. In particular, this is true in the case in which the configuration space is a *Cayley graph*, or simply has a large automorphism group, as FFTs for many finite groups exist. For example, in the case in which the underlying group is abelian, the classical Cooley-Tukey FFT provides a fast algorithm. Other groups with FFTs include symmetric groups and their wreath products, supersolvable groups, and the Lie groups of finite type. See [12] for a survey of recent results.

2. Quasi-Abelian Cayley Graphs

Definition 1. Let G be a finite group and let S be a symmetric set of generators of G , i.e., $\langle S \rangle = G$, $S = S^{-1}$, and $\iota \notin S$, where ι is the identity of G . A graph $\Gamma(G, S)$ with vertex set G and edges $\{s, t\}$ if and only if $t^{-1}s \in S$ is called a Cayley graph. The characteristic function of S will be denoted by $\Theta : G \rightarrow \{0, 1\}$.

Cayley graphs are vertex transitive and hence regular.

Definition 2. A Cayley graph $\Gamma(G, S)$ is called quasi-abelian if S is the union of some conjugacy classes of G .

Clearly, a Cayley graph on a commutative group is quasi-abelian, since each group element forms its own conjugacy class in this case. Some interesting properties of quasi-abelian Cayley graphs are discussed in [13, 14]. Below we shall see that certain algebraic properties of Cayley graphs with abelian groups generalize to quasi-abelian Cayley graphs.

In the case of Cayley graphs we have to distinguish between the “Fourier series expansion” with respect to the graph $\Gamma(G, S)$, equ.(2), and the representation theoretical Fourier transformation on the group G itself. It does not come as a surprise that there is an intimate connection between the two. In fact, the connection

between the algebraic properties of $\Gamma(\mathbf{G}, S)$ and the representation theory of the underlying group \mathbf{G} derives from the following simple facts: The *regular representation* $\boldsymbol{\rho}_{\text{reg}}$ of \mathbf{G} is defined by

$$\boldsymbol{\rho}_{\text{reg}}(s)f(t) = f(s^{-1}t) \quad (5)$$

for any $f : \mathbf{G} \rightarrow \mathbb{C}$. Substituting Θ for f we find $\boldsymbol{\rho}_{\text{reg}}(s)\Theta(t) = \Theta(s^{-1}t) = 1$ if $\{t, s\}$ is an edge of $\Gamma(\mathbf{G}, S)$ and 0 otherwise. Thus we may write the adjacency matrix $\mathbf{A}(\mathbf{G}, S)$ of $\Gamma(\mathbf{G}, S)$ in the form

$$\mathbf{A}(\mathbf{G}, S) = \sum_{s \in S} \boldsymbol{\rho}_{\text{reg}}(s) \quad (6)$$

Definition 3. For any function $f : \mathbf{G} \rightarrow \mathbb{C}$ and any matrix representation $\varrho = \{\boldsymbol{\rho}(s)\}_{s \in \mathbf{G}}$ of \mathbf{G} we call the matrix sum

$$\widehat{f}(\varrho) = \sum_{x \in \mathbf{G}} f(x)\boldsymbol{\rho}(x) \quad (7)$$

the (group theoretic) Fourier Transform of f at ϱ .

Consider a complete set $\{\varrho^1, \dots, \varrho^h\}$ of inequivalent irreducible matrix representations of \mathbf{G} . Let d_k denote the dimension of ϱ^k . Then

$$f(s) = \frac{1}{|\mathbf{G}|} \sum_{k=1}^h d_k \text{Tr} \boldsymbol{\rho}^k(s^{-1}) \widehat{f}(\varrho^k) \quad (8)$$

inverts the Fourier transform.

Following e.g. [15, 8A] we assume that the irreducible representations ϱ^k are unitary, i.e., that $\boldsymbol{\rho}^k(t)^* = \boldsymbol{\rho}^k(t^{-1})$ and introduce

$$\tilde{\rho}_{ij}^k(s) := \sqrt{d_k} \boldsymbol{\rho}_{ji}^k(s^{-1}) \quad (9)$$

These functions are orthonormal w.r.t. the scalar product

$$\langle \varphi | \psi \rangle = \frac{1}{|\mathbf{G}|} \sum_{s \in |\mathbf{G}|} \varphi(s) \psi^*(s) \quad (10)$$

and form a new basis for the vector space of functions of G . Now we are in the position to state the main result of this section.

Theorem 1. Let $\Gamma(\mathbf{G}, S)$ be a quasi-abelian Cayley graph with a finite group \mathbf{G} .

(i) The function $\varepsilon_{ij}^k : \mathbf{G} \rightarrow \mathbb{C}$ defined as

$$\varepsilon_{ij}^k(u) = \frac{1}{\sqrt{|\mathbf{G}|}} \tilde{\rho}_{ij}^k(u) = \sqrt{\frac{d_k}{|\mathbf{G}|}} \boldsymbol{\rho}_{ij}^k(u^{-1}) \quad (11)$$

is an eigenvector of $\mathbf{A}(\Gamma)$ with eigenvalue $\Lambda_k = \frac{1}{d_k} \sum_{s \in S} \chi_k(s)$ where $\chi_k(s) = \text{Tr} \boldsymbol{\rho}^k(s)$ is the character of ϱ^k at s .

(ii) All quasi-abelian Cayley graphs on \mathbf{G} have a common basis of eigenvectors and hence their adjacency matrices commute.

(iii) A function $f : \mathbf{G} \rightarrow \mathbb{R}$ can be expanded in the form

$$f(s) = \sum_{ijk} a_{ij}^k \varepsilon_{ij}^k(s) \quad \text{with} \quad a_{ij}^k = \sqrt{\frac{d_k}{|\mathbf{G}|}} \widehat{f}_{ji}(\varrho^k) \quad (12)$$

Proof. (i) We verify by explicit computation that $\tilde{\rho}_{ij}^k$ is an eigenvector of the adjacency matrix:

$$\begin{aligned}
\sum_{u \in \mathbf{G}} \mathbf{A}_{vu} \tilde{\rho}_{ij}^k(u) &= \sum_{u \in \mathbf{G}} \Theta(vu^{-1}) \tilde{\rho}_{ij}^k(u) \\
&= \sum_{u \in \mathbf{G}} \left\{ \frac{1}{|\mathbf{G}|} \sum_{r,s,t} \sqrt{d_r} \hat{\Theta}_{ts}(\rho^r) \tilde{\rho}_{st}^r(uv^{-1}) \right\} \tilde{\rho}_{ij}^k(u) \\
&= \sum_{u \in \mathbf{G}} \frac{1}{|\mathbf{G}|} \sum_{r,s,t} \hat{\Theta}_{ts}(\varrho^r) \sum_y \tilde{\rho}_{ys}^{r*}(u) \tilde{\rho}_{yt}^r(v) \tilde{\rho}_{ij}^k(u) \\
&= \sum_{r,s,t} \hat{\Theta}_{ts}(\varrho^r) \sum_y \tilde{\rho}_{yt}^r(v) \frac{1}{|\mathbf{G}|} \sum_{u \in \mathbf{G}} \tilde{\rho}_{ij}^k(u) \tilde{\rho}_{ys}^{r*}(u) \\
&= \sum_{r,s,t} \hat{\Theta}_{ts}(\varrho^r) \sum_y \tilde{\rho}_{yt}^r(v) \delta_{kr} \delta_{iy} \delta_{js} = \sum_t \hat{\Theta}_{tj}(\varrho^k) \tilde{\rho}_{it}^k(v)
\end{aligned}$$

Here we have used that $\rho^k(st^{-1}) = \rho^k(s)\rho^k(t^{-1}) = \rho^k(s)\rho^{k*}(t)$ translates to

$$\sqrt{d_r} \tilde{\rho}_{st}^r(vu^{-1}) = \sum_{y=1}^h \tilde{\rho}_{ys}^{r*}(u) \tilde{\rho}_{yt}^r(v)$$

Next we use the fact that Θ is a class function. Hence its Fourier transform is diagonal

$$\hat{\Theta}(\rho^k) = \frac{1}{d_k} \sum_{s \in S} \chi_k(s) \mathbf{I}_{d_k} \quad (13)$$

where $\chi_k(s) = \text{Tr} \rho^k(s)$ is the character of the representation ϱ^k at s . We have therefore

$$\sum_{u \in \mathbf{G}} \mathbf{A}_{vu} \tilde{\rho}_{ij}^k(u) = \sum_t \frac{1}{d_k} \sum_{s \in S} \chi_k(s) \delta_{tj} \tilde{\rho}_{it}^k(v) = \frac{1}{d_k} \sum_{s \in S} \chi_k(s) \times \tilde{\rho}_{ij}^k(v) \quad (14)$$

Changing the normalizations back to the standard scalar product of \mathbb{C} leads to claim (i) of the theorem.

(ii) We have just shown that $\{\tilde{\rho}^{ij}\}$ is an orthonormal basis of eigenvectors of \mathbf{A} whenever S is the union of conjugacy classes of \mathbf{G} . Thus the adjacency matrices of all quasi-abelian Cayley graphs on \mathbf{G} share a common orthonormal basis of eigenvectors. Since the adjacency matrices are symmetric they commute under these circumstances.

(iii) Follows directly from (i), the Fourier inversion formula equ.(8), and the definition of ρ_{ij}^k . \square

Theorem 1 generalizes the following well known result (see e.g. [16]) since all irreducible representations of an abelian group are 1-dimensional.

Corollary 2. *Let \mathbf{G} be a commutative group, and let S be a symmetric set of generators of \mathbf{G} . Then the irreducible characters χ_k of \mathbf{G} are eigenvectors of $\mathbf{A}(\mathbf{G}, S)$ with corresponding eigenvalue $\Lambda_k = \sum_{s \in S} \chi_k(s)$.*

3. Hamming graphs and Fitness landscapes

3.1. Hamming graphs. The vertex set of a Hamming graph \mathcal{Q}_α^n consists of all sequences of length n composed from an alphabet of size α . Two vertices x and y are joined by an edge iff they differ in exactly one sequence position, i.e., if there is exactly one $k \in \{1, \dots, n\}$ such that $x_k \neq y_k$ and $x_i = y_i$. It is not hard to verify that \mathcal{Q}_α^n arises as a Cayley graph of the group $Z_\alpha^n = (\mathbb{Z}/\alpha\mathbb{Z})^n$ with a set of generators of the form $e_k^l = (\iota_1, \dots, \iota_{k-1}, e_k^l, \iota_{k+1}, \dots, \iota_n)$, where $e_k^l \neq \iota_k$ and ι_k denotes the identity on the k -th copy of $(\mathbb{Z}/\alpha\mathbb{Z})$. Corollary 2 yields the eigenvalues of eigenvectors of Hamming graphs [17, 18]. The eigenfunctions are known as *Walsh functions* (see eg. [12, 19])

$$\varepsilon_I = \exp\left(\frac{2\pi i}{\alpha} \langle x, I \rangle\right), \quad (15)$$

the corresponding eigenvalues are $\Lambda_p = p\alpha$ with multiplicities $(\alpha - 1)^p \binom{n}{p}$, where $p = |I|$ is the ‘‘interaction order’’.

3.2. FFTs for Hamming graphs. For the efficient evaluation of Fourier transforms on groups of the form $(\mathbb{Z}/\alpha\mathbb{Z})^n$, where α and n are positive integers, we appeal to the *separation of variables* technique (see eg. [12]). This technique can best be illustrated by the following example.

Let f be any complex-valued function $f : (\mathbb{Z}/\alpha\mathbb{Z})^2 \rightarrow \mathbb{C}$. We wish to calculate the Fourier transform of f . That is, we wish to evaluate

$$\hat{f}(\alpha_1, \alpha_2) = \sum_{m_1=0, m_2=0}^{n-1} \omega^{m_1\alpha_1} \omega^{m_2\alpha_2} f(m_1, m_2), \quad (16)$$

where $\omega = e^{-2\pi i/\alpha}$ and $0 \leq \alpha_1, \alpha_2 < \alpha$. For a given α_1, α_2 , evaluating the above sum would take α^2 operations. Therefore, to compute the discrete Fourier transform (DFT) of f via this naive algorithm would take α^4 operations.

We can improve on this operation count by first rewriting the above equation as a multiple sum:

$$\begin{aligned} \hat{f}(\alpha_1, \alpha_2) &= \sum_{m_2=0}^{\alpha-1} \omega^{m_2\alpha_2} \sum_{m_1=0}^{\alpha-1} \omega^{m_1\alpha_1} f(m_1, m_2) \\ &= \sum_{m_2=0}^{\alpha-1} \omega^{m_2\alpha_2} \tilde{f}(\alpha_1, m_2) \end{aligned}$$

Next, we evaluate $\tilde{f}(\alpha_1, m_2)$ (i.e. the inner sum) for the α many different values of m_2 . That is to say, we compute α many 1-dimensional DFTs of length α , one DFT for each m_2 . This can be done in at most α^3 operations. (We allow for the possibility that the individual, $1 - d$ DFTs are evaluated by direct sum, without use of the FFT.) After doing this, the outer sum

$$\sum_{m_2=0}^{n-1} \omega^{m_2\alpha_2} \tilde{f}(\alpha_1, m_2)$$

is evaluated for the α -many different values of α_1 . In other words, there are α -many DFTs of length α to perform, this time one for each α_1 . As above, these DFTs can all be accomplished in at most α^3 operations. Therefore, by rewriting the sum

in equ.(16) as a multiple sum, the Fourier transform of f can now be computed in $2\alpha^3$ operations.

By introducing multiple sums, we effectively arranged the input data $f(m_1, m_2)$ as a square array of dimensions $\alpha \times \alpha$. Doing the separate sums can be interpreted as first taking the DFT of all the rows of the array, followed by taking the DFT of all the columns. This notion of replacing a single sum with multiple sums, one for each dimension, is the basic idea behind the separation of variables technique.

Note that α^{2n} operations are required to compute the Fourier transform of a function f defined on $(\mathbb{Z}/\alpha\mathbb{Z})^n$ via the naive algorithm. The separation of variables technique enables us to compute the Fourier transform in only $n\alpha^{n+1}$ operations. We remind the reader that depending on how the individual DFTs are performed, computing the Fourier transform of f in fewer operations is possible.

To actually perform the evaluations, we use an algorithm given in [20]. Philosophically, the Fourier transform of a function f on $(\mathbb{Z}/\alpha\mathbb{Z})^n$ is accomplished as follows. The data is arranged in a contiguous block of memory of length α^n . Consecutive blocks of length α are interpreted as the rows of a $\alpha^{n-1} \times \alpha$ matrix \mathbf{X} .

1. Overwrite each length α row of \mathbf{X} (i.e. the consecutive length α blocks) with that row's DFT.
2. Replace \mathbf{X} with its transpose.
3. Repeat the first two steps $n - 1$ more times.

After n iterations, the array \mathbf{X} will contain the Fourier transform of f .

3.3. RNA Landscapes. In this section we consider a class typical landscapes from computational biology. Ribonucleic acid (RNA) molecules are linear polymers formed from four different monomers, guanine, cytosine, adenosine, and uracil, designated by the letters $\{\mathbf{G}, \mathbf{C}, \mathbf{A}, \mathbf{U}\}$. These nucleotides form three types of stable base pairs, \mathbf{GC} , \mathbf{AU} , \mathbf{GU} (in the order of decreasing strength) which cause the linear polymer to fold back on itself, see Figure 1. The actual three-dimensional structure of an RNA molecule can be approximated by its *secondary structure*, that is, the list of the base pairs that minimize the structure's energy. The dominant role of the secondary structure is well documented in nature by the conservation of secondary structure elements in evolution [21].

There are efficient dynamic programming algorithms that allows the prediction of the energetically most favored RNA secondary structure given the sequence of nucleotides [22, 23]. These algorithm are polynomial in both run time and memory requirements. Hence they allow a detailed analysis of fitness landscapes derived from RNA structures, see e.g., [24] and the references therein. The computations of the RNA folding landscapes reported below are performed using the **Vienna RNA Package**, release 1.2 (Nov. 1997). This software is freely available on the internet¹.

A first attempt at computing approximate amplitude spectra for RNA landscape were reported in [25] based on estimating the correlation function $r(s)$ from random walks. The usefulness of complete Fourier transforms of RNA landscapes is discussed in [10]. Here we show that the application of FFT techniques allows the investigations of the generic features of RNA landscapes in regime of sequence length that it no longer plagued with the strong finite size effects that make the interpretation very difficult when the sequences are so short that most of them do not form a stable structure.

¹<http://www.tbi.univie.ac.at/>




Figure 1. The folding of the one-dimensional primary structure (sequence) of an RNA molecule is decomposed conceptually into two steps: (i) Folding of the sequence into a planar secondary structure by formation of complementary Watson-Crick base pairs **GC** and **AU** as well as the weaker **GU** pairs. (ii) Formation of the 3D tertiary structure from the planar pattern. The reason for this decomposition is that the intramolecular forces stabilizing the secondary structures — base pairing and base pair stacking — are much stronger than those accounting for the particular spatial arrangement. Thus the free energy for the formation of the entire spatial structure can be estimated by the free energy for the formation of the secondary structure. While the computation of 3D structures from the pure sequence information is impossible at present, the secondary structure can be readily calculated.

In figure 2 we compare four different RNA folding landscapes: the energy landscape (assigning the ground state energy to each RNA sequence), the ensemble free energy landscape, and two landscapes based on structure distances to two different targets. The restriction to pure **GC** sequences allows us to study longer sequences in which finite size effects are less important. While the four landscapes are based on quite different principles, energy of folding and structural similarity to a fixed anchor point, respectively, we observe that their amplitude spectra are very similar. This is a consequence of the fact that all four landscapes can be understood (at least approximately in the case of the folding energies) as a composition of the genotype-phenotype map (mapping the sequence onto its secondary structure) and a (close to linear) evaluation of the secondary structure [26]. As a consequence, all four landscapes mostly reflect the regularities of the underlying genotype-phenotype map.

For instance, the strong distinction between even and odd modes reflects the dominating effect of base pair stacking, which always involves an even number of nucleotide positions. Similarly, the increase of the contributions of larger values p can be attributed to the increase in average stack (helix) length. The average stack length eventually settles down to a constant value [27].

The same generic features can be seen in the amplitude spectra computed for the true four-letter alphabet **GCAU** as shown in Figure 3. The main difference is the virtual absence of the linear $p = 1$ mode in the **GC** landscape, while the additive

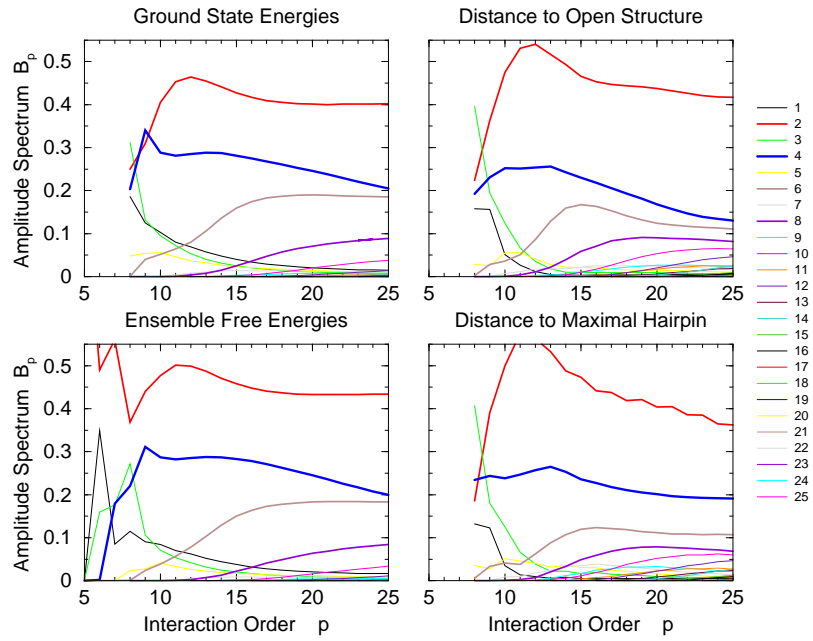


Figure 2. Amplitude spectra of RNA folding landscapes with **GC** alphabet as a function of the chain length n .

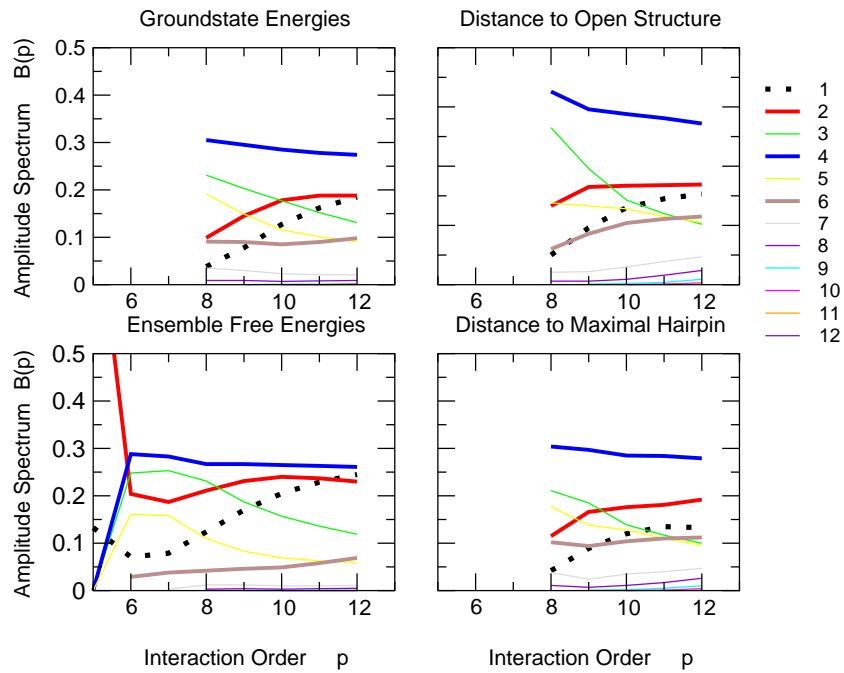


Figure 3. Amplitude spectra of RNA folding landscapes with **GCAU** alphabet as a function of the chain length n .

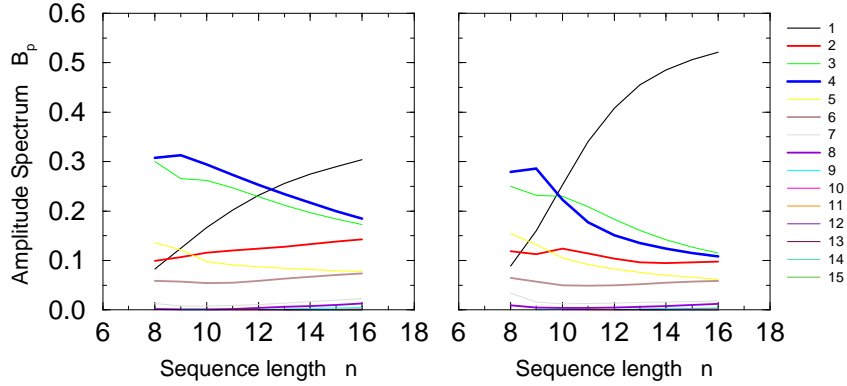


Figure 4. Amplitude spectra of RNA folding energy landscapes with **GCA** and **GCU** alphabet, respectively. Note the strong difference in the amplitude of the linear mode.

mode accounts for a substantial fraction of the variance in the **GCAU** case. This is due to the different average strength of **GC** and **AU** base pairs, which introduced a substantial dependence of both the folding energy and structural composition on the **G+C** content of the sequence.

Artificially restricted alphabets can be useful to retrieve information on the overall structure of RNA landscapes. In Figure 4 we compare the **GCA** and **GCU** alphabets. Their main difference lies in the base pairing logic: in the former case **A** is always unpaired, resulting in an increase of folding energy with the **A**-content. In the **GCU** case, there is the possibility to form **GU** pairs, thus reducing the effect of the **U**-concentration. However, the amplitude of the linear mode increases since **C-U** substitutions have a much smaller influence on the structure – and hence act predominantly linearly.

4. Neutral Network Landscapes

Sometimes only the distinction between viable and non-viable phenotypes is important. In this case a landscape reduces to a Boolean function $f : V \rightarrow \{0, 1\}$, which we shall call a *neutral network landscape*. Examples arise from RNA folding, where we may consider only sequences folding into a particular secondary structure [28, 29], in evolving CA landscapes [30], and in random graph models such as [31, 32]. In many cases neutral network landscapes are restricted to a (sometimes small) linear subspace of the sequence space.

Let f be a neutral network landscape on Z_2^n , with non-zero values restricted to the m -dimensional linear subspace $Z_2^m < Z_2^n$. Furthermore, suppose the distribution of 0s and 1s in Z_2^m is i.i.d., with a probability q for encountering a 1. For short we call this model a “homogeneous neutral network with density q ”.

Theorem 3. *The expected amplitude spectrum of a homogeneous neutral network of density q confined to an m -dimensional subspace of Z_2^n is*

$$B_p = q 2^{-(n-m)} \binom{n-m}{p} + (1-q) 2^{-n} \binom{n}{p} \quad (17)$$

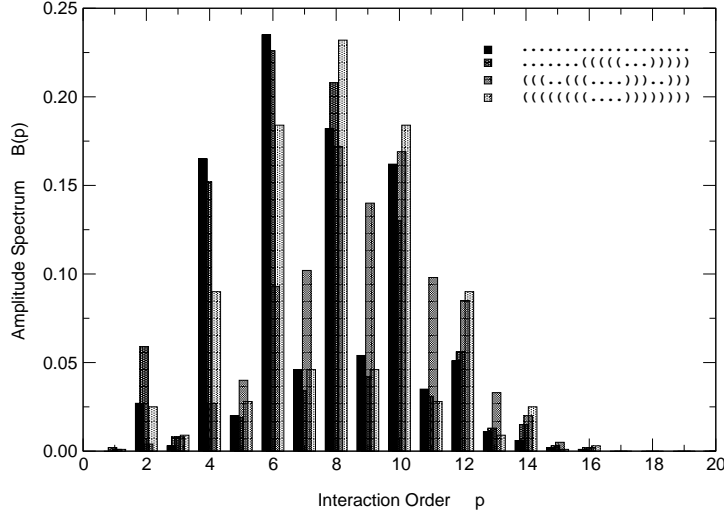


Figure 5. Amplitude spectra of four RNA neutral networks for the **GC** alphabet and chain length $n = 20$. Typical networks, such as those of the open structure (black), the most frequent structure (dark gray), and the maximally paired structure (light gray) are very similar. Again, the difference between even and odd interaction orders is clearly visible. For neutral network of the marginally stable network of the structure $((\dots((\dots))\dots))$ we observe a much smaller difference between even and odd modes. This reflects the fact that almost all point mutations change this particular structure.

Proof. Since Suppose $f(x) = 0$ in $x \notin Z_2^m$ we have

$$\hat{f}(y) = \sum_x (-1)^{xy} f(x) = \sum_{x \in Z_2^m} (-1)^{xy^*} f(x) \tag{18}$$

where y^* denotes the projection of y onto Z_2^m . Thus $\hat{f}(y) = \hat{f}(y')$ whenever $y^* = y'^*$.

In order to evaluate the expected amplitude spectrum of f , we need to evaluate $\mathbb{E}[\hat{f}(y)]$ for $y \in Z_2^m$, since

$$\mathbb{E}[\beta_p] = \sum_{y: |y|_1=p} \mathbb{E}[|\hat{f}(y^*)|^2] = \sum_{j=0}^p \sum_{\substack{y \in Z_2^m \\ |y|_1=j}} \mathbb{E}[\hat{f}(y^*)^2] \binom{n-m}{p-j} \tag{19}$$

The last binomial coefficient accounts for all y^* that have $p - j$ '1's in the "tail". A simple computation shows that

$$\begin{aligned} \mathbb{E}[|\hat{f}(y^*)|^2] &= \sum_{x,z} (-1)^{yx} (-1)^{yz} \mathbb{E}[f(x)f(z)] \\ &= q^2 \sum_{x \neq z} (-1)^{yx} (-1)^{yz} + q \sum_x ((-1)^{xy})^2 \\ &= q^2 \sum_x \sum_{x \neq z} (-1)^{yx} \sum_z (-1)^{yz} + q(1-q)2^m \\ &= q^2 4^m \delta_{y,0} + q(1-q)2^m \end{aligned} \tag{20}$$

Thus we find

$$\begin{aligned}\mathbb{E}[\beta_p] &= q^2 4^m \binom{n-m}{p} + \sum_{j=0}^m q(1-q) 2^m \binom{m}{j} \binom{n-m}{p-j} \\ &= q^2 4^m \binom{n-m}{p} + q(1-q) 2^m \binom{n}{p}\end{aligned}\tag{21}$$

The normalized amplitude spectrum is obtained by dividing this expression by $\sum_p \mathbb{E}[\beta_p] = q^2 2^{n+m} + q(1-q) 2^{n+m} = q 2^{n+m}$ \square

Amplitude spectra of RNA landscapes are of course more complicated than this idealized model. In particular, the difference between pairs and unpaired positions is reflected by the amplitude difference between odd and even interaction orders. The over-all shape of the amplitude spectra in Figure 5, which resemble binomial distributions (overlaid by the even/odd distinction) with a mean somewhere between 7 and 8, on the other hand, nicely reflect the restriction of neutral networks to lower-dimensional sub-spaces of \mathcal{Q}_2^n .

5. Permutation Problems

5.1. The Representation Theory of the Symmetric Group. The irreducible representations of the symmetric group on n symbols, S_n are in a natural one-to-one correspondence with the nondecreasing partitions of n . Any such partition $\lambda = (\lambda_1, \dots, \lambda_h)$ (with $\sum_i \lambda_i = n$, $\lambda_i \geq \lambda_{i+1} > 0$) of n (denoted $\lambda \vdash n$) corresponds to a permutation action of S_n on the space of Young tableaux of shape λ . Within this (in general) reducible action exists a well-defined irreducible constituent, thereby giving the associated irreducible representation.

As several of the following examples reduce to consideration of some specific instances of these permutation representations on Young tableaux we briefly recall here their construction.

Let $\lambda \vdash n$. A *Young tableau* (named for the British mathematician Alfred Young) is constructed by filling in a *Ferrers* or *Young diagram* of shape λ with the numbers $1, \dots, n$ without repetition. The Young diagram of shape λ is a left-justified array of square boxes with λ_i boxes in row i .



$$\tag{22}$$

If the numbers $1, \dots, n$ are placed without repetition in the boxes of a Young diagram of shape λ , then a *Young tableau of shape λ* is defined. Two Young tableaux of shape λ are said to be (row) equivalent if they differ only by permuting the entries within a given row. An equivalence class of Young tableaux of a fixed shape is called a *tabloid* of the same shape. A given tabloid is denoted by forming a representative Young tableau, and removing the internal vertical lines. Equation (23) shows three

equivalent tableaux of shape $(4, 3, 1, 1)$ and their associated tabloid.

$$\begin{array}{|c|c|c|c|} \hline 3 & 1 & 4 & 2 \\ \hline 7 & 6 & 5 & \\ \hline 8 & & & \\ \hline 9 & & & \\ \hline \end{array}
 \quad
 \begin{array}{|c|c|c|c|} \hline 2 & 4 & 3 & 1 \\ \hline 6 & 7 & 5 & \\ \hline 8 & & & \\ \hline 9 & & & \\ \hline \end{array}
 \quad
 \begin{array}{|c|c|c|c|} \hline 1 & 2 & 3 & 4 \\ \hline 6 & 7 & 5 & \\ \hline 8 & & & \\ \hline 9 & & & \\ \hline \end{array}
 \quad
 \begin{array}{|c|c|c|c|} \hline 1 & 2 & 3 & 4 \\ \hline 5 & 6 & 7 & \\ \hline 8 & & & \\ \hline 9 & & & \\ \hline \end{array}
 \tag{23}$$

Let X^λ denote the set of tabloids of shape λ . The symmetric group on n symbols, S_n , acts transitively on X^λ by permuting the entries in the tabloids. The subgroup stabilizing a given tabloid is called a *Young subgroup*. Note that for fixed λ , each such subgroup is isomorphic to $S_{\lambda_1} \times \cdots \times S_{\lambda_h}$. Let S_λ denote the particular Young subgroup which permutes the sets $\{1, \dots, \lambda_i\}, \dots, \{\lambda_1 + \cdots + \lambda_{h-1} + 1, \dots, n\}$ within themselves. This gives rise to an identification between X^λ and S_n/S_λ and a permutation representation of S_n on the vector space of complex-valued functions on X^λ . The permutation module is denoted M^λ .

The representation theory of S_n is studied by decomposing in a systematic way the permutation representations M^λ . For present purposes it suffices to say that within each M^λ there is a uniquely determined irreducible subspace S^λ , and letting λ run through all partitions of n accounts for all the irreducible representations of S_n , without multiplicity. Sagan gives a very nice exposition of the construction [33]. James and Kerber provide an encyclopedic account with many references [34].

The decomposition of the permutation module M^λ immediately gives a decomposition of the vector space of functions on S_n defined by the associated matrix elements. It is the decomposition of this space of functions which is of interest for the analysis of the optimization problems. To obtain some feeling for this connection it is perhaps easiest to work within the context of a specific example, so consider the representation $M^{(n-1,1)}$, given by the symmetric group S_n acting on Young tableaux of shape $(n-1, 1)$. Any such Young tableau is determined by the entry in the second row, and thus, may be identified with the action of S_n on the standard basis e_1, \dots, e_n given by $\rho(\pi)e_i = e_{\pi(i)}$, which is the so-called ‘‘defining representation’’ of S_n . For any fixed i , the matrix elements $\{\rho_{i1}, \dots, \rho_{in}\}$ span an S_n -invariant subspace of $L(S_n)$. This follows from the fact that $(\tau\rho_{ij})(\sigma) = \rho_{ij}(\tau^{-1}\sigma)$, implying that $\tau\rho_{ij} = \rho_{i\tau(j)}$. Thus, the set of matrix elements $\{\rho_{i1}, \dots, \rho_{in}\}$ do themselves span a copy of $M^{(n-1,1)}$, thereby providing n easily identified isomorphic copies of the space $M^{(n-1,1)}$. The representation space $M^{(n-1,1)}$ decomposes as

$$M^{(n-1,1)} = S^{(n)} \oplus S^{(n-1,1)}$$

where $S^{(n)}$ denotes the trivial representation, spanned by the subspace of vectors with constant coordinates, while $S^{(n-1,1)}$ denotes its $n-1$ -dimensional orthogonal irreducible complement of those vectors whose coordinates sum to zero.

These copies of $M^{(n-1,1)}$ are not mutually orthogonal. For example, each contains the same copy of the trivial representation. Also, notice that the values of any one row of matrix elements, $\{\rho_{i1}, \dots, \rho_{in}\}$ is determined by knowing all the other rows. In fact, the matrix elements span a space of dimension $(n-1)^2 + 1$ which as a representation of S_n has an irreducible decomposition isomorphic to $S^{(n)} \oplus (n-1)S^{(n-1,1)}$ and for any function $f : S_n \rightarrow \mathbb{C}$ the computation of $\hat{f}(\rho)$ is equivalent to computing the projection of f onto the trivial representation, as well as the isotypic component of $L(S_n)$,

$$L(S_n) = \oplus_{\lambda \vdash n} I^\lambda \tag{24}$$

Table 1. The lowest Laplacian eigenvalues of $\Gamma(\mathbb{S}_n, \tau)$.

$\lambda \vdash n$	d_λ	χ_λ/d_λ	Λ_λ	Remark
(n)		1	0	constant
$(n-1, 1)$		$1 - 2/(n-1)$	n	LAP, equ.(26)
$(n-2, 2)$		$1 - 4/n$	$2(n-1)$	symmetric TSP
$(n-2, 1, 1)$		$1 - 4/(n-1)$	$2n$	“antisymmetric” TSP
$(n-3, 3)$		$1 - 6/n - 6/(n(n-1))$	$3(n-2)$?

which corresponds to the irreducible representation $S^{(n-1,1)}$, denoted as $I^{(n-1,1)}$ in (24). These projections are the Fourier transforms at the corresponding irreducible representations and in this sense, the projection onto $I^{(n-1,1)}$ encodes the pure first order information about f . More generally, for each of the M^λ , an irreducible decomposition has a corresponding block diagonalization of the representation, with repeated blocks, and the irredundant choice of matrix elements corresponds to an isotypic decomposition of the space of matrix elements.

5.2. Transpositions and Cayley Graphs of \mathbb{S}_n . For most combinatorial optimization problems defined on permutations, transpositions seem to be at least a natural move set. Reversals, or canonical transpositions are other choices that frequently appear in practice.

For simplicity we will restrict ourselves here to the Cayley graphs $\Gamma(\mathbb{S}_n, \tau)$ in which neighborhood is given by transpositions. By construction it is quasi-abelian, and hence the discussion in section 2 applies.

The Laplacian eigenvalues can be obtained in this case directly from the Frobenius equation

$$\frac{\chi_{(\lambda_1, \lambda_2, \dots, \lambda_h)}(\tau)}{d_{(\lambda_1, \lambda_2, \dots, \lambda_h)}} = \frac{1}{n(n-1)} \sum_{j=1}^h [\lambda_j^2 - (2j-1)\lambda_j] \quad (25)$$

The lowest Laplacian eigenvalues are listed in table 1.

5.3. The Linear Sum Assignment Problem. An assignment problem consists of n facilities and n locations and a cost c_{ik} of establishing facility i in location k . In additions, there will be costs involved in transporting goods from one location to the other and flows of goods that depend on the facilities. In the simplest case the cost function is of the form

$$f(\pi) = \sum_{i=1}^n c_{i, \pi(i)} \quad (26)$$

where (c_{ij}) is an arbitrary $n \times n$ matrix. We shall show below that only the representations (n) and $(n-1, 1)$ contribute to f . In other words, f is elementary and belongs to the eigenvalue $\Lambda_{(n-1,1)}$ when considered as a landscape on the Cayley graph $\Gamma(\mathbb{S}_n, \tau)$.

Let us define the functions ρ_{ij} (for $1 \leq i, j \leq n$) on \mathbb{S}_n by

$$\rho_{ij}(\pi) = \delta_{j, \pi(i)} = \begin{cases} 1 & \text{if } \pi(i) = j \\ 0 & \text{otherwise.} \end{cases} \quad (27)$$

With the usual action of \mathbb{S}_n on the ρ_{ij} we have

$$(\sigma \rho_{ij})(\pi) = \rho_{ij}(\sigma^{-1}\pi) = \rho_{i, \sigma(j)}(\pi) \quad (28)$$

and thus

$$\sigma \rho_{ij} = \rho_{i,\sigma(j)}. \quad (29)$$

Notice that in fact the ρ_{ij} are the matrix elements for the “standard representation” of S_n as permutation matrices (i.e., $\rho_{ij}(\pi)$ is the i, j entry of the permutation matrix associated to π acting on the standard basis, e_1, \dots, e_n).

It is easy to see that each of the sets of matrix elements $\{\rho_{i1}, \dots, \rho_{in}\}$ spans an invariant subspace V_i of the full vector space spanned by all of the ρ_{ij} , which we denote by V . Furthermore, V_i is isomorphic to the permutation module $M^{(n-1,1)}$ given by the action on the Young tableaux of shape $(n-1, 1)$, which has as its irreducible decomposition

$$M^{(n-1,1)} \cong S^{(n)} \oplus S^{(n-1,1)}$$

where $S^{(n)}$ is the trivial representation given by the subspace of constant functions, and $S^{(n-1,1)}$ its orthogonal complement of vectors whose coefficients (in the standard basis) add to zero.

Notice that the V_i are not mutually orthogonal. For example, the values $\rho_{11}, \dots, \rho_{1n}, \dots, \rho_{n-1,1}, \dots, \rho_{n-1,n}$ at π determine the values $\rho_{ni}(\pi)$, $i = 1, \dots, n$. It is not too hard to check that

$$V \cong (n-1)S^{(n-1,1)} \oplus S^{(n)}. \quad (30)$$

The relation with the linear problem is now immediate. Defining $f \in V$ by

$$f = \sum_{ij} c_{ij} \rho_{ij} \quad (31)$$

and substituting equ.(27) we obtain the general form of the cost function equ.(26). Thus, since $f \in V$, we will only have contributions from the $(n-1, 1)$ -isotypic, and the constant subspace.

The LAP assignment problem can be solved by a variety of techniques in polynomial time. For a recent review see [35].

5.4. Quadratic Assignment Problems. In the simplest class of hard assignment problems there is not only a cost c_{ik} for establishing facility i in location k but also a flow f_{ij} between facility i and facility j , a distance (or cost rate per unit flow) d_{kl} between locations k and l , and a cost c_{ij} for setting up facility i at location j . The total cost of a particular assignment $\pi \in S_n$ is then

$$f(\pi) = \sum_{i=1}^n \sum_{j=1}^n f_{ij} d_{\pi(i)\pi(j)} + \sum_{i=1}^n c_{i,\pi(i)} \quad (32)$$

The *quadratic assignment problem* (QAP) is then to minimize $f(\pi)$. For a survey see [36]. The QAP contains the Traveling Salesman Problem (TSP) as a special case: setting $c_{ij} = 0$ and $f_{ij} = \delta_{i,j-1}$ (with indices interpreted mod n) yields $f(\pi) = \sum_{j=1}^n d_{\pi(j)\pi(i+1)}$.

In the following we show that only the representations (n) , $(n-1, 1)$, and the two “quadratic modes” $(n-2, 2)$ and $(n-2, 1, 1)$ contribute to the function (32). We can neglect the linear term $\sum_i c_{i,\pi(i)}$ here since we have already shown above that it contains only components of (n) and $(n-1, 1)$.

We start with the functions

$$\rho_{(i,j),(k,l)}(\sigma) = \begin{cases} 1 & \text{if } (\sigma(i), \sigma(j)) = (k, l) \\ 0 & \text{otherwise.} \end{cases} \quad (33)$$

for $i \neq j$. We observe that the functions $\rho_{(i,j),(k,l)}$ are the matrix elements for the permutation representation of S_n acting on the Young tableaux of shape $(n-2, 1, 1)$, denoted $M^{(n-2,1,1)}$. This has an irreducible decomposition

$$M^{(n-2,1,1)} \cong S^{(n)} \oplus 2S^{(n-1,1)} \oplus S^{(n-2,2)} \oplus S^{(n-2,1,1)}.$$

Note that this is **not** the irreducible decomposition of the vector space V spanned by the functions $\rho_{(i,j),(k,l)}$, but V has contributions only from the irreducibles above. If we define the function $f \in V$ by

$$f = \sum_{i,j,k,l} f_{ij} d_{k,l} \rho_{(i,j),(k,l)} \quad (34)$$

then by (33)

$$f(\pi) = \sum_{ij} f_{ij} d_{\pi(i), \pi(j)} \quad (35)$$

and thus should only have nonzero contributions from the partitions (n) , $(n-1, 1)$, $(n-2, 2)$, $(n-2, 1, 1)$.

If the distance matrix D is symmetric, $d_{ij} = d_{ji}$, we consider the symmetrized functions

$$\eta_{(i,j),\{k,l\}} = (\rho_{(i,j),(k,l)} + \rho_{(i,j),(l,k)}). \quad (36)$$

Then,

$$\eta_{(i,j),\{k,l\}}(\sigma) = \begin{cases} 1 & \text{if } \{\sigma(i), \sigma(j)\} = \{k, l\} \\ 0 & \text{otherwise.} \end{cases} \quad (37)$$

Thus, the $\eta_{(i,j),\{k,l\}}$ are the matrix coefficients for the permutation representation $M^{(n-2,2)}$ and thus span a subrepresentation within $M^{(n-2,1,1)}$. This implies that if we let W denote the subspace spanned by the $\eta_{(i,j),\{k,l\}}$ then $V = W \oplus W^\perp$ and that W^\perp is a direct sum of the $(n-2, 1, 1)$ -isotypic and a piece isomorphic to some number of copies of $(n-1, 1)$.

If D is antisymmetric, $d_{ij} = -d_{ji}$, then the projection of f onto the $\eta_{(i,j),\{k,l\}}$ is 0, so that in the antisymmetric case there is no $(n-2, 2)$ component. Conversely, this shows that in the symmetric case with $d_{ij} = d_{ji}$, there can be no projection onto the $(n-2, 1, 1)$ -isotypic.

The classification in Table 1 hence generalized to QAPs in the following way: If D is symmetric, then f is built from (n) , $(n-1, 1)$, and $(n-2, 2)$ components. If D is anti-symmetric, only (n) , $(n-1, 1)$, and $(n-2, 1, 1)$ contribute. We remark that in general the projection onto the $(n-1, 1)$ -isotypic does not vanish even if $C = 0$. It would be interesting to see whether the cost functions with vanishing $(n-1, 1)$ -component correspond to an interesting subclass of QAP.

5.5. Higher order examples. While each of these optimization problems contain information only in a subset of the irreducibles there are more general instances in which the relevant data analysis can require the full set of Fourier transforms. Data on the symmetric group can occur in the form of “ranked data”, in which a population is asked to rank n items in order of preference. For example, Diaconis

has used this approach to analyze the results of the presidential election for the American Psychological Association [37]. We give only a very brief synopsis here.

As part of the voting process, association members are given a ballot sheet and asked to rank the candidates (labeled (1) – (5)) in order of increasing preference. In this way, each member picks a permutation of the 5 candidate labels. This stage in the vote thus yields a function $f : S_5 \rightarrow \mathbb{C}$ in which $f(\pi)$ is the number of voters with preference order π . The Fourier transforms at the irreducible representations of S_5 naturally encode “coalition” information, corresponding to the associated partition. For example, the Fourier transform at $(3, 2)$ reflects the popularity of the various pairs of candidates; the Fourier transform at $(2, 2, 1)$ can be interpreted as the popularity of a particular threesome, but then considered as ordered within the threesome as a pair and a singleton. Other data sets are considered in [15, 38].

Acknowledgments. D.N.R. is supported by SFI and NSF Presidential Faculty Fellowship NSF DMS-9553134, under which P.K. is supported. Partial support for P.F.S. in connection with his research visit to SFI in Sep. 1999 by a grant from the Austrian Federal Ministry for Science and Transportation is gratefully acknowledged. A major part of this work was performed during visits to the Santa Fe Institute in 1999 where it received support from grants to SFI from the John D. and Catherine T. MacArthur foundation, NSF, DOE, and the Keck Foundation.

References

- [1] R. Grone, R. Merris, and V. Sunder. The Laplacian spectrum of a graph. *SIAM J. Matrix Anal. Appl.*, 11:218–238, 1990.
- [2] B. Mohar. The Laplacian spectrum of graphs. In Y. Alavi, G. Chartrand, O. R. Ollermann, and A. J. Schwenk, editors, *Graph Theory, Combinatorics, and Applications*, pages 871–898, New York, 1991. John Wiley and Sons, Inc.
- [3] B. Mohar. Some applications of Laplace eigenvalues of graphs. In G. Hahn and G. Sabidussi, editors, *Graph Symmetry: Algebraic Methods and Applications*, volume 497 of *NATO ASI Series C*, pages 227–275. Kluwer, Dordrecht, 1997.
- [4] F. R. K. Chung. *Spectral Graph Theory*, volume 92 of *CBMS*. American Mathematical Society, 1997.
- [5] Y. C. de Verdière. Multiplicités des valeurs propres Laplaciens discrete at laplaciens continus. *Rendiconti di Matematica*, 13:433–460, 1993.
- [6] J. Friedman. Some geometric aspects of graphs and their eigenfunctions. *Duke Math. J.*, 69:487–525, 1993.
- [7] L. K. Grover. Local search and the local structure of NP-complete problems. *Oper. Res. Lett.*, 12:235–243, 1992.
- [8] P. F. Stadler. Linear operators on correlated landscapes. *J. Physique I (France)*, 4:681–696, 1994.
- [9] P. F. Stadler. Landscapes and their correlation functions. *J. Math. Chem.*, 20:1–45, 1996.
- [10] W. Hordijk and P. F. Stadler. Amplitude spectra of fitness landscapes. *Adv. Complex Systems*, 1:39–66, 1998.
- [11] P. F. Stadler and G. Tinhofer. Equitable partitions, coherent algebras and random walks: Applications to the correlation structure of landscapes. *MATCH*, 40:215–261, 1999.
- [12] D. K. Maslen and D. N. Rockmore. Generalized FFTs – a survey of some recent results. In L. Finkelstein and W. Kantor, editors, *Proceedings of the DIMACS Workshop on Groups and Computation, June 7-10, 1995*, pages 183–237. DIMACS, 1997.
- [13] J. Wang and M. Xu. Quasi-abelian Cayley graphs and Parsons graphs. *Eur. J. Comb.*, 18:597–600, 1997.
- [14] B. Zgrablic. On adjacency-transitive graphs. *Discrete Math.*, 182:321–332, 1998.
- [15] P. Diaconis. *Group Representations in Probability and Statistics*. Inst. of Math. Stat., Hayward, CA, 1989.
- [16] L. Lovász. Spectra of graphs with transitive groups. *Periodica Math. Hung.*, 6:191–195, 1975.

- [17] D. Rumschitzky. Spectral properties of Eigen's evolution matrices. *J. Math. Biol.*, 24:667–680, 1987.
- [18] A. W. M. Dress and D. S. Rumschitzki. Evolution on sequence space and tensor products of representation spaces. *Acta Appl. Math.*, 11:103–111, 1988.
- [19] N. J. Fine. The generalized walsh functions. *Trans. Am. Math. Soc.*, 69:66–77, 1950.
- [20] C. Van Loan. *Computational Framework for the Fast Fourier Transform*. SIAM, Philadelphia, 1992.
- [21] T. R. Cech. Conserved sequences and structures of group I introns: building an active site for RNA catalysis — a review. *Gene*, 73:259–271, 1988.
- [22] M. Zuker and D. Sankoff. RNA secondary structures and their prediction. *Bull. Math. Biol.*, 46:591–621, 1984.
- [23] I. L. Hofacker, W. Fontana, P. F. Stadler, S. Bonhoeffer, M. Tacker, and P. Schuster. Fast folding and comparison of RNA secondary structures. *Monatsh. Chemie*, 125:167–188, 1994.
- [24] P. Schuster, P. F. Stadler, and A. Renner. RNA structures and folding: From conventional to new issues in structure predictions. *Curr. Opinions Structural Biol.*, 7:229–235, 1997.
- [25] R. Happel and P. F. Stadler. Canonical approximation of fitness landscapes. *Complexity*, 2:53–58, 1996.
- [26] P. F. Stadler. Fitness landscapes arising from the sequence-structure maps of biopolymers. *J. Mol. Struct. (THEOCHEM)*, 463:7–19, 1999.
- [27] W. Fontana, D. A. M. Konings, P. F. Stadler, and P. Schuster. Statistics of rna secondary structures. *Biochemistry*, 33:1389–1404, 1993.
- [28] W. Grüner, R. Giegerich, D. Strothmann, C. M. Reidys, J. Weber, I. L. Hofacker, P. F. Stadler, and P. Schuster. Analysis of RNA sequence structure maps by exhaustive enumeration. II. Structures of neutral networks and shape space covering. *Monath. Chem.*, 127:375–389, 1996.
- [29] C. M. Reidys, P. F. Stadler, and P. Schuster. Generic properties of combinatory maps: Neural networks of RNA secondary structures. *Bull. Math. Biol.*, 59:339–397, 1997.
- [30] W. Hordijk. Correlation analysis of the synchronizing-CA landscape. *Physica D*, 107:255–264, 1997.
- [31] S. Gavrillets and J. Gravner. Percolation on the fitness hypercube and the evolution of reproductive isolation. *J. Theor. Biol.*, 184:51–64, 1997.
- [32] C. M. Reidys. Random induced subgraphs of generalized n -cubes. *Adv. Appl. Math.*, 19:360–377, 1997.
- [33] B. Sagan. *Representation theory of the symmetric group*. Wadsworth, New York, 1992.
- [34] G. D. James and A. Kerber. *The representation theory of the symmetric group*, volume 16 of *Encyclopedia of Mathematics*. Addison-Wesley, Reading, MA, 1981.
- [35] R. E. Burkard and E. Çela. Linear assignment problems and extensions. In Z. Du and P. Pardalos, editors, *Handbook of Combinatorial Optimization*, page to appear. Kluwer Academic Publishers, Dordrecht, 1999. Univ. Graz and Technical Univ. Graz (Austria); SFB F003 report #127.
- [36] P. Pardalos, F. Rendl, and H. Wolkowicz. The quadratic assignment problem: A survey and recent developments. In *Proceedings of the DIMACS Workshop on Quadratic Assignment Problems*, volume 16 of *DIMACS Series in Discrete Mathematics and Theoretical Computer Science*, pages 1–41. American Mathematical Society, 1994.
- [37] P. Diaconis. A generalization of spectral analysis with applications to ranked data. *Ann. Stat.*, 17:949–979, 1989.
- [38] D. N. Rockmore. Some applications of generalized FFTs. In L. Finkelstein and W. Kantor, editors, *Proceedings of the DIMACS Workshop on Groups and Computation, June 7-10, 1995*, pages 329–369, 1997. An appendix with D. Healy.