Shapes of tree representations of spin-glass landscapes

Wim Hordijk[a], **José F. Fontanari**[a]§, **and Peter F. Stadler**[b,c,d]

[a]Instituto de Física de São Carlos, Universidade de São Paulo, Caixa Postal 369, 13560-970 São Carlos SP, Brazil

[b]Bioinformatik, Institut für Informatik. Universität Leipzig, Kreuzstraße 7b, D-04103 Leipzig, Germany

[c]Institut für Theoretische Chemie und Molekulare Strukturbiologie, Universität Wien, Währingerstraße 17, A-1090 Wien, Austria

[d]The Santa Fe Institute, 1399 Hyde Park Road, Santa Fe, NM 87501, USA

Abstract.

Much of the information about the multi-valley structure of disordered spin systems can be convened in a simple tree structure – a barrier tree – the leaves and internal nodes of which represent, respectively, the local minima and the lowest energy saddles connecting those minima. Here we apply several statistics used in the study of phylogenetic trees to barrier trees that result from the energy landscapes of p-spin models. These statistics give information about the shape of these barrier trees, in particular about balance and symmetry. We then ask if they can be used to classify different types of landscapes, compare them with results obtained from random trees, and investigate the structure of subtrees of the barrier trees. We conclude that at least one of the used statistics is capable of distinguishing different types of landscapes, that the barrier trees from p-spin energy landscapes are quite different from random trees, and that subtrees of barrier trees do not reflect the overall tree structure, but their structure is correlated with their "depth" in the tree.

PACS numbers: 75.10.Nr, 87.23.Kg

§ To whom correspondence should be addressed.

1. Introduction

The notion of energy (fitness) landscapes has played a crucial role in the development of many areas of physics and biology such as disordered systems, neural networks, combinatorial optimization problems [1], RNA folding [2], and evolutionary change [3, 4], to mention only a few. In particular, considerable effort has been devoted to the study of the interplay between the geometry of the landscape and the nature of the relaxation dynamics, searching heuristic or evolutionary process unfolding on the landscape. However, the inherent high-dimensionality of these landscapes poses a serious hindrance to the characterization of their topology. In fact, most of the studies have focused on the statistical characterization of a few local properties of the landscape by looking at, e.g., the auto-correlation function of unbiased walks over the configuration space [5] or the energy distribution of local minima [6, 7], while a satisfactory description of a landscape should address also the (relative) energy differences of the local minima, the height of the barriers between these minima, as well as the distribution of saddle-points [8, 9].

The idea of condensing all the landscape information into a tree structure, termed *barrier tree*, was introduced in the context of RNA and protein folding [10, 11, 12, 13], and spin-glass models [8, 14, 15, 16]. The advantage of barrier trees, whose leaves represent the local minima and the internal nodes the lowest-energy saddles connecting those minima, is that they are both visually appealing (much information can be obtained from just looking at them, see figure 1), as well as mathematically well-defined, lending themselves to rigorous analysis [17]. However, a general quantitative measure to characterize unambiguously different kinds of barrier trees (and hence energy landscapes) remains to be obtained. For instance, the size-frequency distribution of low energy saddles $\psi(w) \sim w^{-D}$, where $w = w(s)$ is the fraction of minima that can be connected through saddle s, does not provide a good measure because $D \sim 2$ regardless of whether the barrier tree results from a spin-glass landscape or whether it is generated randomly [16].

The situation seems to be similar in the analysis of phylogenetic trees, where it is also believed that the shape of the tree contains valuable clues about the evolutionary process [18, 19], but so far no *single* satisfactory measure of tree-shape has been constructed [20, 21, 22, 23, 24]. In this contribution we apply five measures of tree shape that were originally used to study phylogenetic trees (see, e.g., [22]) to barrier trees resulting from the Ising p-spin model. These measures provide statistical information about the shape of the barrier tree, mainly its symmetry or balance, but ignore the branching lengths, i.e. the height of the barriers between minima. While the extreme statistics of these heights provide useful information on the performance of local search algorithms such as simulated annealing [15], a measure based solely on the shape of the barrier tree seems more adequate to classification purposes since the shape is probably insensitive to variations in minor details of the underlying energy landscape. Tree shapes are important in the biological context also because many methods of phylogeny estimation, including parsimony, do not produce branch lengths. We find that all five measures can be used to tell random from spin-glass trees, but only one measure can distinguish between different spin-glass trees.

In the next section, a brief overview of the Ising p-spin model and energy landscapes is given. Section 3 then reviews the notion of barrier trees. In section 4, phylogenetic trees and several tree shape statistics used to study them are discussed. In particular, we point out that the subtree that connects two leaves in a barrier tree corresponds to the evolutionary path of minimum fitness cost, in contrast with the traditional phylogenetic trees for which the subtrees are determined by the similarity between the leaves, regardless of the barrier height between them. The results of applying these measures to barrier trees of p-spin energy landscapes are presented in section 5. The main conclusions are then summarized in the final section of the paper.

2. Energy landscapes of p-spin models

Consider a system of N Ising spins $s = (s_1, \ldots, s_N)$ where $s \in \{-1, +1\}$, with the following energy function

$$H_p(s) = -\sum_{1 \leq i_1 < i_2 \ldots < i_p \leq N} J_{i_1 i_2 \ldots i_p} s_{i_1} s_{i_2} \ldots s_{i_p}. \tag{1}$$

Here, $1 \leq p \leq N$ and the $J_{i_1 i_2 \ldots i_p}$ are i.i.d. random variables from a Gaussian distribution with mean 0 and variance $p!/(2N^{p-1})$. Thus, each of the 2^N possible spin configurations is assigned an energy value that is completely specified by the (magnetic) interactions within all possible subsets of p spins. This spin glass is known as the p-spin model [25]. The limit $p \to \infty$ reduces to the random energy model (REM) where each spin configuration s gets an independent random energy value from a Gaussian distribution of mean 0 and variance N [25]. Another important limiting case is $p = 2$, which corresponds to the Sherrington-Kirkpatrick (SK) model [26].

We call two spin configurations s and t *neighbors* if they differ in only one of the N spins, i.e., when $\sum_{i=1}^{N} |s_i - t_i| = 2$. In other words, neighboring configurations can be turned into each other by a single spin flip. When p is small, neighboring spin configurations will have highly correlated energy values, since the one spin s_i that is different between the two configurations only influences a small number of the possible subsets of p spins. When p increases, the energy values of neighboring spin configurations will become less correlated, becoming completely uncorrelated in the limit $p \to \infty$.

The energy landscape of a p-spin model consists of the configuration space V of the 2^N possible spin configurations, with the single spin flip neighborhood relation imposed on it, and where the energy value $H_p(s)$ of each spin configuration s is considered to be its "height". This gives rise to the intuitive image of a more or less mountainous landscape with peaks, valleys, and saddle points. A "walk" on this landscape consists of moving from one neighboring spin configuration to another, climbing up a peak or going down a valley, or perhaps just moving around randomly. A local search algorithm such as simulated annealing can be seen as performing such a walk, in search of the lowest

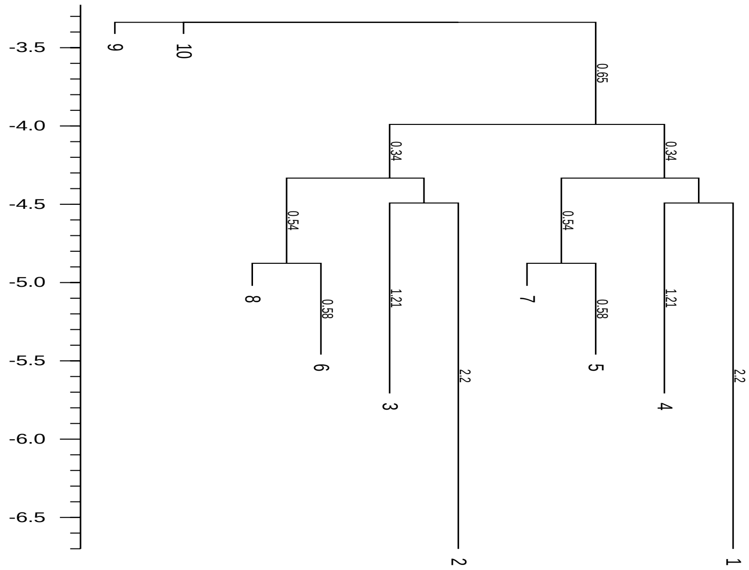

Figure 1. Barrier tree for $N = 10$ and $p = 2$. The local minima are labeled 1 to 10, and the height of the barriers is shown along the branches. The actual energy values can be read from the scale on the left.

valley. In addition, p-spin landscapes have been used to model evolutionary processes (see, e.g., [27, 28, 29]) since they form a class of tunably rugged landscapes similar to Kauffman's Nk-model [3]. In this context, evolution is described as an "adaptive" walk on the energy landscape.

3. Barrier trees

A *local minimum* in a p-spin energy landscape is simply a spin configuration s that has a lower energy than all of its neighbors. A *path* $\vec{p}_{st}$ between two configurations s and t is a sequence of neighboring configurations, starting at s and ending at t. In other words, it represents a series of single spin flips that transforms configuration s into t. Note that there exist multiple paths between any pair of configurations s and t. A *saddle point* between two local minima s and t is then defined as the minimum from the set of maximum energy values along each possible path $\vec{p}_{st}$ between s and t; see e.g. [8, 9]. So, the energy value $E[s,t]$ of this saddle point is

$$E[s,t] = \min_{\vec{p}_{st}} \left\{ \max_{z \in \vec{p}_{st}} \{H_p(z)\} \right\}. \tag{2}$$

The *barrier* $B(s)$ of a local minimum s is defined as the height of the lowest saddle point that connects s with a local minimum t of lower energy,

$$B(s) = \min_{t} \{E[s,t] - H_p(s) | H_p(t) < H_p(s)\}. \tag{3}$$

The information about the energy values of a landscape's local minima and the barriers that connect these local minima can be represented by a *barrier tree*. In such a tree, the leaves of the tree represent the local minima, and the internal nodes represent the barriers, with branches between local minima and the barriers that connect them. Figure 1 shows an example of a barrier tree for an $N = 10$ and $p = 2$ p-spin landscape. There are 10 local minima in this landscape (labeled 1 to 10 in the tree), with 9 saddle points (the internal nodes). The length of each branch in the tree indicates the height of the corresponding barrier (this value is shown along the branch).

The algorithm for constructing these barrier trees is presented in [13, 17]. It is implemented in the `barriers` program‖, which constructs the tree from a sorted list of energy values of all spin configurations in the landscape. The program `barriers` is used here to generate barrier trees of p-spin landscapes, see also [15, 16] for applications to spin glass problems.

4. Phylogenetic trees

Phylogenetic trees are often used to study the historical relations within or between groups of biological species; see e.g. [19]. The currently existing species (or subspecies) form the leaves of the tree, and two related species are linked through their last common ancestor, which forms an internal node in the tree. The length of the branch between a species and its ancestor indicates how long ago the speciation event occurred that led to the current species. "Dead-ends" in this tree represent extinction events. So, the shape of a phylogenetic tree contains information about patterns of speciation and extinction (and possibly about the rates of these occurrences), and thus tells us something about the evolution of different species.

The traditional approaches to reconstruct or infer a phylogenetic tree given the extant species (i.e., the leaves) are based on sequence similarity, i.e., configurational overlap only, which can be justified by the (usually implicit) assumptions of a flat fitness landscape and a diffusive behavior in sequence space [19]. However, it seems intuitive that, regardless of the similarity between two sequences, if the fitness costs of all possible evolutionary paths connecting them are high, they must be put far apart in the phylogenetic tree. In the barrier tree approach both similarity and fitness cost are taken into account to yield the evolutionary path of *minimum fitness cost* connecting two species. Thus it may be viewed as a generalization of the maximum parsimony principle to rugged fitness landscapes. We recall that maximum parsimony chooses the tree (or trees) that require the fewest evolutionary changes (spin flips, in the present context). In figure 2 we give an example in which the distance-based tree and barrier tree have different topologies.

Various methods have been proposed to analyze phylogenetic trees. Here, we consider five statistics that were used in [22] to measure tree symmetry and balance. The trees are assumed to be binary trees with n leaves (or species) and thus $n - 1$ internal nodes, with the root being the last common ancestor of all n species. Let $d(i, j)$ be the graph-theoretical distance between two nodes of the tree, i.e., the number of

‖ The source code is available at `http://www.tbi.univie.ac.at/~ivo/RNA/Barriers/`.

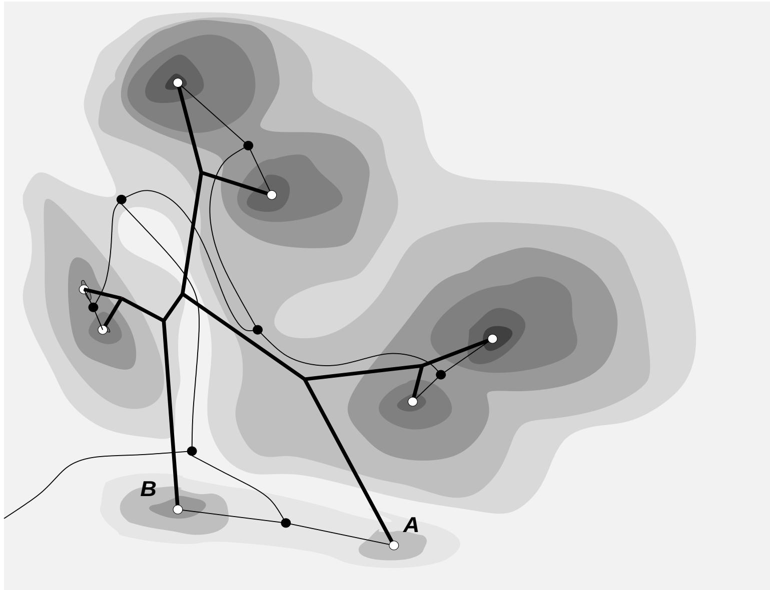

Figure 2. The tree obtained from the minimum fitness paths, shown by thin lines with the saddle points indicated by black dots and the leaves by white dots, does not necessarily coincide with the trees obtained from clustering methods that are based on sequence similarity, shown here with thick lines. The darker regions indicate higher fitness configurations.

edges along the path that connects them. Furthermore, we denote the root of the tree by $\varnothing$. The *height* of a leaf l is $h_l = d(\varnothing, l)$. Equivalently, h_l is the number of internal nodes between leaf l and the root $\varnothing$ (inclusive). For each interior node i we have two subtrees with r_i and s_i leaves, respectively. We assume $r_i \geq s_i$. The *subtree-height* of an interior node i is $m_i = \max_{l \in T_i} d(i, l)$ where the maximum is taken over all leaves l in the subtree T_i below i, i.e., the subtree of which i is the root.

With this notation we may define the following five characteristic values for the shape of a binary rooted tree:

(i) $H = \frac{1}{n} \sum_{l=1}^{n} h_l$ is the *average height* of a leaf in the tree¶.

(ii) $\sigma_H = \sqrt{\frac{1}{n} \sum_{l=1}^{n} (h_l - H)^2}$, is the standard deviation of the leaf height.

(iii) $C = \frac{2}{n(n-3)+2} \sum_{i=1}^{n-1} (r_i - s_i)$ is a measure for the *imbalance* of the tree. Up to normalization it is the same as Colless's imbalance measure [20]. A closely related measure of this type, which we will not use here, essentially amounts to averaging $(r_i - s_i)/(r_i + s_i)$ instead, see e.g. [23, 24].

(iv) $B_1 = \sum_{i \neq \varnothing} 1/m_i$ is the average inverse subtree height, where the sum is taken over all $n-2$ internal nodes i excluding the root $\varnothing$.

(v) $B_2 = \sum_{l=1}^{n} 2^{-h_l} h_l$ is an alternatively weighted average leaf-height.

In [22] the variance σ_H^2 was considered, but here we will use the standard deviation. Both H and σ_H have larger values for more asymmetric trees. In [22] it is shown that

¶ In [22] this quantity is denoted by $\bar{N}$.

the expected value of $\langle H \rangle = 2\sum_{k=2}^{n} 1/k$ for random trees with n leaves. Values of H larger than this indicate trees more asymmetric than a random tree. For a completely symmetric tree, $\sigma_H = 0$, while it has a maximum value for a completely asymmetric tree.

The imbalance measure C examines the internal nodes of a tree. It "weighs" the subtrees branching out from each internal node by counting and comparing the number of leaves in each subtree. These weight differences are then averaged and normalized over all internal nodes of the tree. The value of C increases from 0 for a completely symmetric tree to 1 for a completely asymmetric tree. The quantity B_1 looks at the longest possible path m_i from each internal node i to any of the leaves in its subtree. The statistic B_2 is based on an index of information content. For highly asymmetric trees B_2 will quickly converge to a value of 2. For a completely symmetric tree, is will be equal to $\log_2(n)$, where n is the number of leaves in the tree. Both B_1 and B_2 have smaller values for increasingly asymmetric trees.

5. Results

Here, we apply the statistics presented in the previous section to barrier trees that result from p-spin energy landscapes. The parameter values used are $N = 10$, 12, 15, 18, 20, and 22, and $p = 2$, 3, and ∞ (REM). For each combination of N and p, 100 independent landscapes were generated randomly, and the barrier trees for each of these landscapes were constructed. The tree statistics reported here are the averages of the 100 trees for each parameter combination.

Three of the statistics (H, σ_H, and B_1) are exponential in N, the number of spins. For example, the relation between H and N is $H = 0.43\exp(0.20N)$ for $p = 2$. The data for H (symbols) and the exponential fits (solid lines) are shown in a semi-log plot in figure 3a for the different values of p. The data for σ_H and B_1 (not shown) are similar, but with different slopes. It turns out that the average number of leaves, n, in the barrier trees is also exponential in the number of spins N, with the same slope as for H. Analytical values for this exponent, defined as $\lim_{N\to\infty} \ln n/N$, for the p-spin models are derived e.g. in [6] (see also [30]). The numerical values from the simulations reported here are compiled in table 1.

The imbalance C is sub-linear in N, as shown in figure 3b, and eventually converges to 1 for large N (so the logarithm fitting cannot remain valid for all N). However, the value of B_2 very quickly converges to 2 with increasing N, as shown in Figure 3c, and thus it does not seem to be a very useful measure to distinguish the trees. It should be noted at this point that the number n of leaves, i.e., local minima, depends not only on N but also on the type of the landscape. We will therefore consider tree measures as a function of n rather than N below. In fact, n is the natural parameter in the analysis of phylogenetic trees, obtained by simply counting the number of leaves in the tree.

The data presented so far explicitly assumed that we already know the values of N and p of the energy landscapes. Suppose we do not know these values, but all we

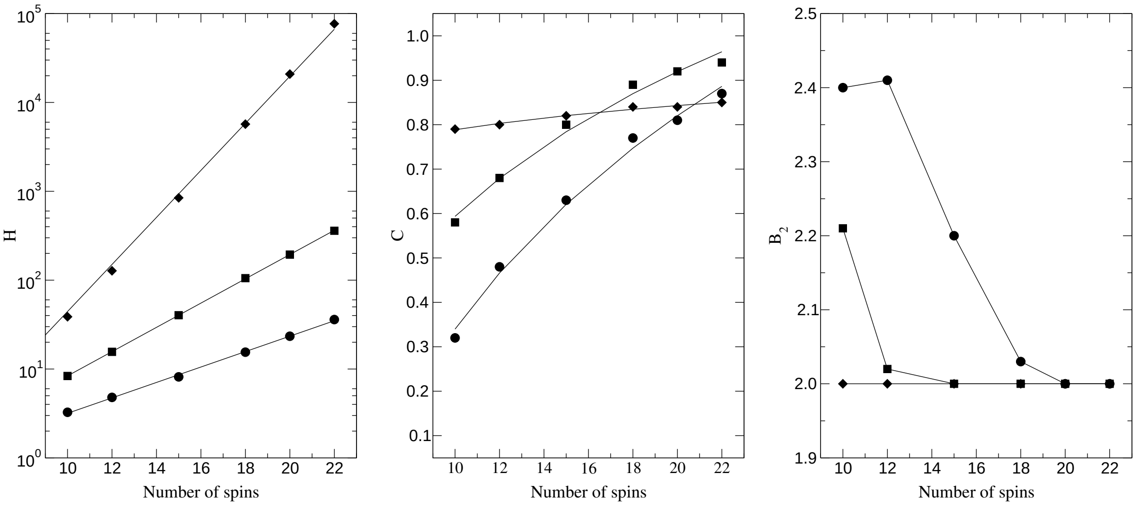

(a) (b) (c)

Figure 3. Size dependence of tree statistics as a function of the number of spins N for three classes of landscapes: p-spin models with $p = 2$ (●) and $p = 3$ (■), and the REM (♦). Each data point is an average over 100 independent landscapes. The solid line is an exponential fit to the data for the average leaf height H in (a). The solid lines in (b) are fits to $C = a + b \ln N$. We do not have a justification for this particular functional form. The values of B_2 quickly converge to the asymptotic value of 2 in (c), where the solid lines are simply guides to the eye.

have are several instances (say, 100) of barrier trees that were derived from some p-spin landscape. Is it possible to say something about the underlying landscape from just the barrier tree? (The problem faced by the phylogeneticists is even harder since they have access to a single tree of arbitrary size only [22].)

In figure 4a, the values of H of all landscapes are lumped together and plotted against n, the (average) number of leaves in the tree, in a log-log plot. This results in a straight line, indicating a power-law behavior. In fact, the slope in this case is 1, since the slopes of H and n when plotted against N are equal (see table 1). The plots for σ_H and B_1 (not shown) are similar, but with slopes slightly larger and slightly smaller than 1, respectively.

In this case, it is not possible to distinguish the trees that result from, e.g., the

	$p = 2$	$p = 3$	REM
H	0.1999	0.3157	0.6340
σ_H	0.2485	0.3381	0.6374
B_1	0.1070	0.1832	0.6040
n	0.1951	0.2945	0.6306

Table 1. The slopes of the exponential fits for the statistics H, σ_H, B_1, and n, for each value of p.

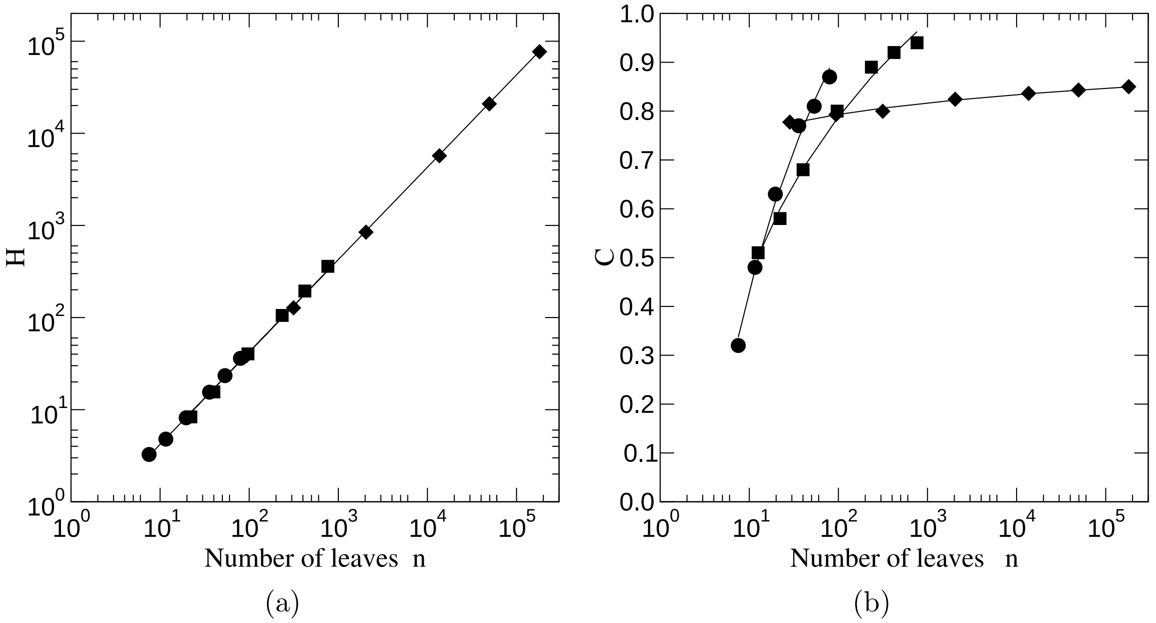

Figure 4. (a) The values of H and (b) the values of C plotted against n, the average number of leaves in the tree. We find that H does not depend on the p for different p-spin models. On the other hand, the values of C show a significant dependence on the details of the landscape. The fit shown is $C = a + b \ln\ln n$ and the convention is • $p = 2$, ■ $p = 3$, and ♦ REM.

$p = 2$ and the $p = 3$ landscapes: they all fall on the same line. Since B_2 converges to a fixed value for larger trees, it is not very useful for distinguishing different trees either. However, when the imbalance statistic C is plotted against n, there is a clear distinction, as can be seen in figure 4b. Plotting C against $\ln\ln n$ (n is on a double logarithmic scale), the data can be fitted by straight lines with different slopes for the different values of p. So, C clearly can be used as a statistic to distinguish and classify different barrier trees, and thus their underlying landscapes.

Next we address the question: How different are barrier trees of p-spin landscapes from random trees? To answer it, we generated random binary trees with $n = 10^i$ leaves for $i = 1, 2, 3, 4$, and 5. For each value of i, 100 random trees were generated and the same five statistics were calculated. The random trees were generated as follows. First, create n nodes (the leaves) and put them in a set A. Next, remove two random nodes x and y from A, create a new node z and make x and y its two children, and put z in the set A. Repeat this procedure until there is only one node left in A, which will be the root of the tree.

Figure 5a shows the data for H against n (the number of leaves) in a semi-log plot for random trees. Clearly, H depends logarithmically on n. The results for σ_H and B_2 (not shown) are similar, but with different slopes. Figure 5b shows the data for C against n on a log-log plot. The fit to the data is a power law with exponent -0.81. The results for B_1 (not shown) are similar but with a (positive) exponent very close to 1.

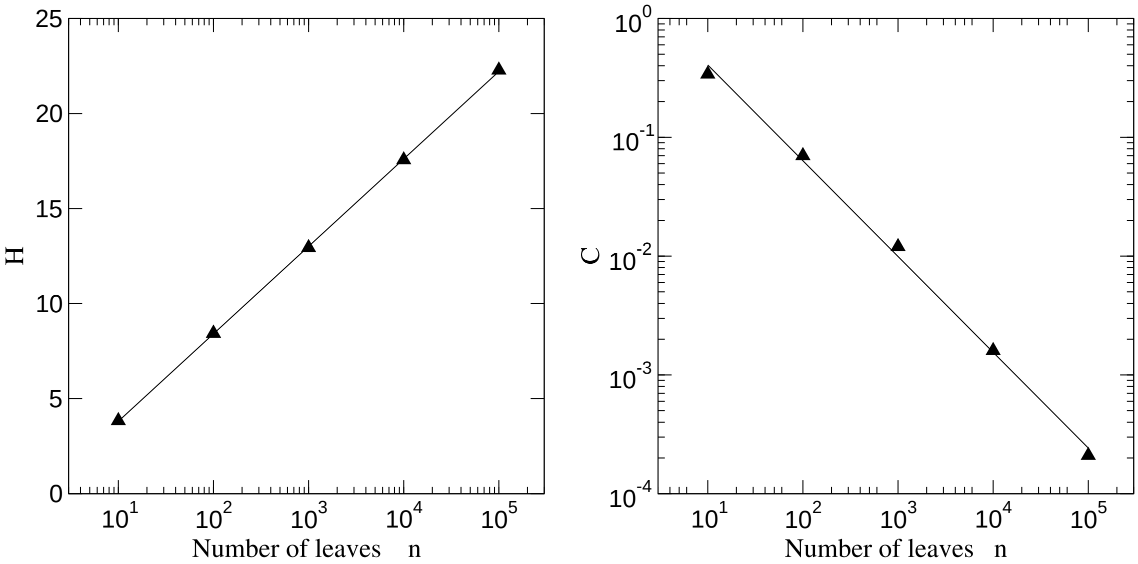

(a) (b)

Figure 5. Average leaf-height H and imbalance measure C as a function of tree-size n for random trees.

The results for random trees are quite different from those for the barriers trees of p-spin landscapes. In the p-spin case, H and σ_H depend exponentially on n, whereas for random trees this dependence is logarithmic. For C, the reverse is true. Moreover, for random trees, C *decreases* with increasing n, while it increases for p-spin barrier trees. For B_1, in both cases the dependence is exponential, but with different exponents. Finally, for B_2 both cases are completely different, with a (downward) convergence to the value $B_2 = 2$ for p-spin trees, and a logarithmic increase for random trees. Clearly, barrier trees from p-spin landscapes are much more asymmetric than random trees. A similar, though qualitative, conclusion was reached by considering the size-distribution of minima connected through a high-energy (i.e., closer to the root) saddle-point [16].

Finally, we investigate the structure of subtrees of barrier trees. Instead of calculating the tree statistics on the entire tree, they are calculated on subtrees starting at some internal node of the tree. This way, it can be determined whether the tree has a self-similar structure or not (i.e., whether subtrees look similar to the tree as a whole). Three different instances of an $N = 20, p = 3$ p-spin landscape were taken and their barrier trees were constructed. Next, the five statistics were calculated on the subtrees starting from each of the $n-1$ internal nodes of these trees (a binary tree with n leaves has $n-1$ internal nodes). Each internal node in a barrier tree represents a saddle point in the underlying energy landscape, and the internal nodes are characterized by the energy values of their corresponding saddle points. Obviously, nodes higher up in the tree (closer to the root) will have higher energy values than nodes lower down (or deeper) in the tree (closer to the leaves).

Figure 6a shows the results for H. The statistics for the subtrees are plotted against

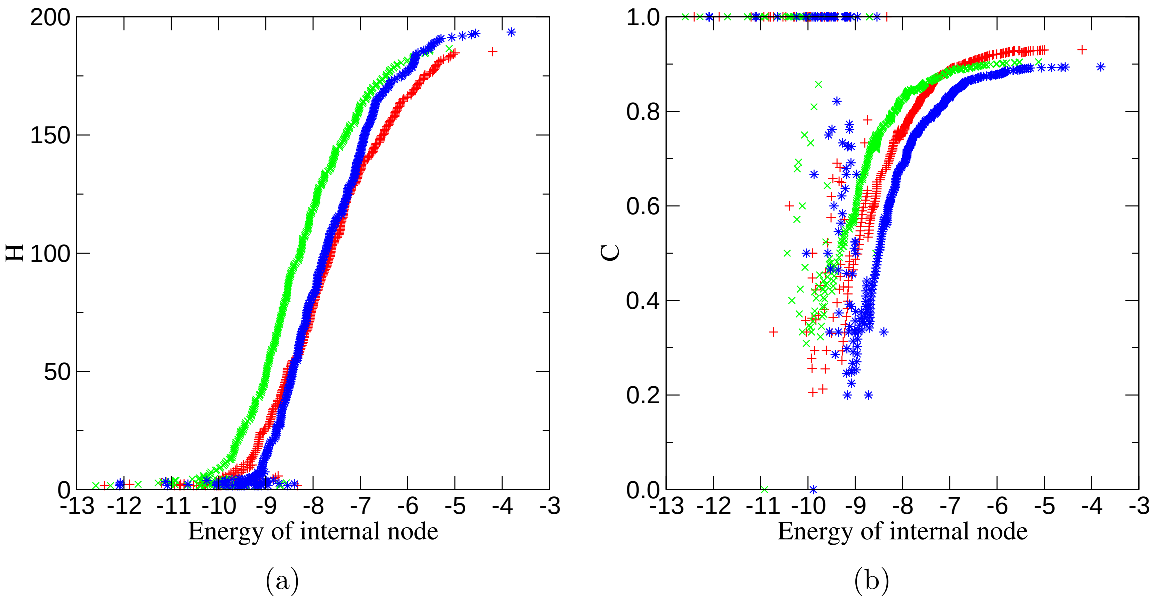

Figure 6. (a) H and (b) C against the energy of the internal node (saddle point) for three different landscapes for $N = 20$ and $p = 3$.

the energy value of the internal node that forms the root of the subtree. The data for the three different landscapes can be distinguished from the slightly different range in energy values, but the overall shape is the same for all three. Clearly, the values for H depend strongly on the energy value of the root of the subtree. The results for σ_H and B_1 (not shown) are similar. Figure 6b shows the results for the statistic C. We find that subtrees with a high-energy root are extremely unbalanced. The same is true for the subtrees with very few leaves that we find near the global optimum. In the intermediate regime we find nodes with very balanced subtrees. For B_2, the values vary much more widely, but with most of the points falling on or near the $B_2 = 2$ line. The overall results for $p = 2$ and for the REM are very similar.

¿From these plots it is clear that the structure of subtrees does not reflect the structure of the tree as a whole. Instead, there is a rather well-defined dependence on the structure of a subtree and its *depth* in the tree, i.e., the energy value of the internal node that forms its root. The landscape structure around local minima and saddle points with a relatively high energy value is therefore significantly different from the structure around local minima and saddle points with a low energy. Since this structure is correlated with the energy value of the local minima and saddle points, this information could possibly be used to guide local search algorithms.

6. Conclusions

One of the main outcomes of the mean-field (replica) theory of spin glasses was the prediction that in the low temperature phase, the phase space is broken into infinitely many pure states or valleys s^a [1]. Searching for universal, in the sense of model independent, features of the distribution of valleys is one of the main goals of the theory [31, 32]. However, the replica theory provides a characterization of the valley structure for nonzero temperature only: at $T = 0$ the overlap between any two pure states, say s^a and s^b, become unity and all fine details of the hierarchy of valleys are lost. To investigate the valley structure in this regime one has then to resort either to the annealed estimates of the correlations between local minima in the thermodynamic limit [6, 7] or to the exact numerical calculation of the barrier trees for relatively small system sizes [8, 9, 16]. As pointed out before, the latter seems a more convenient approach for the purpose of classifying families of landscapes or spin-glass models. Furthermore, the barrier tree also contains information about the structure of subspaces in the landscape around local minima and saddle points, since the structure and symmetry of the subtrees are clearly correlated with their depth in the tree (or the energy level of the node that forms the root of the subtree). This information can probably be used to guide local search algorithms such as simulated annealing in their search for the lowest (or global) minimum.

In searching for efficient measures to characterize landscapes, it was recently shown that the size-frequency distribution of the number of leaves w connected by a saddle-point s is too robust to be a useful measure, as it yields the same power-law $\psi(w) \sim w^{-2}$ for both p-spin barrier trees and random trees [16]. In this contribution, we improve considerably the landscape systematics by considering five measures proposed originally to characterize the shape of phylogenetic trees [22]. Three of the measures used, namely, H, σ_H and B_1 were proved independent of the underlying spin-glass landscape when plotted against the number of leaves n, as evidenced by the "data collapse" illustrated in figure 4a. Only one of the measures, the imbalance C, can be used to differentiate between, e.g., $p = 2$-type and $p = 3$-type energy landscapes. It is remarkable, however, that all five measures yield completely different results for random trees, owing to a different scaling with n. In particular, barrier trees generated from energy landscapes of p-spin models appear to be much more asymmetric than random trees, and the asymmetry increases with increasing N (the number of spins) or n (the number of leaves). It remains a challenge to find (if it exists) a disordered spin system whose associated barrier trees exhibit balance and symmetry properties akin to those of the random trees or, perhaps an easier task, that violate the scaling law shown in figure 4a.

Acknowledgments

Thanks to Christoph Flamm at the University of Vienna, Austria, for his help with the `barriers` program. This research was supported by Fundação de Amparo à Pesquisa do Estado de São Paulo (FAPESP), project 99/09644-9. The work of J.F.F. is supported in part by CNPq and WH is supported by FAPESP.

References

[1] Mézard M, Parisi G and Virasoro M A 1987 *Spin Glass Theory and Beyond* (Singapore: World Scientific)
[2] Schuster P and Stadler P F 1994 *Comput. Chem.* **18** 295
[3] Kauffman S A and Levin S 1987 *J. Theor. Biol.* **128** 11
[4] Kauffman S A 1993 *The Origins of Order* (Oxford: Oxford University Press)
[5] Weinberger E D 1990 *Biol. Cybern.* **63** 325
[6] Gross D J and Mézard M 1984 *Nucl. Phys. B* **240** 431
[7] de Oliveira V M and Fontanari J F 1997 *J. Phys. A: Math. Gen.* **30** 8445
[8] Nemoto K 1988 *J. Phys. A: Math. Gen.* **21** L287
[9] Vertechi A M and Virasoro M A 1989 *J. Phys. France* **50** 2325
[10] Becker O M and Karplus M 1997 *J. Chem. Phys.* **106** 1495
[11] Wales D J, Miller M A and Walsh T R 1998 *Nature* **394** 758
[12] Garstecki P, Hoang T X and Cieplak M 1999 *Phys. Rev. E* **60** 3219
[13] Flamm C, Fontana W, Hofacker I L and Schuster P 2000 *RNA* **6** 325
[14] Klotz T and Kobe S 1994 *J. Phys. A: Math. Gen.* **27** L95
[15] Ferreira F F, Fontanari J F and Stadler P F 2000 *J. Phys. A: Math. Gen.* **33** 8635
[16] Fontanari J F and Stadler P F 2002 *J. Phys. A: Math. Gen.* **35** 1509
[17] Flamm C, Hofacker I L, Stadler P F and Wolfinger M T 2002 *Z. Phys. Chem.* **216** 155
[18] Mooers A O and Heard S B 1997 *Quart. Rev. Biol.* **72** 31
[19] Felsenstein J 2002 *Inferring Phylogenies* (Sunderland: Sinauer Associates)
[20] Colless D H 1982 *Syst. Zool.* **31** 100
[21] Shao K-T and Sokal R R 1990 *Syst. Zool.* **39** 266
[22] Kirkpatrick M and Slatkin M 1993 *Evolution* **47** 1171
[23] Fusco G and Cronk Q C B 1995 *J. Theor. Biol.* **175** 235
[24] Purvis A, Katzourakis A and Agapow P-M 2002 *J. Theor. Biol.* **214** 99
[25] Derrida B 1981 *Phys. Rev. B* **24** 2613
[26] Sherrington D and Kirkpatrick S 1975 *Phys. Rev. Lett.* **35** 1792
[27] Anderson P W 1983 *Proc. Natl. Acad. Sci. USA* **80** 3386
[28] Rokhsar D S, Anderson P W and Stein D L 1986 *J. Mol Evol.* **23** 119
[29] Amitrano C, Peliti L and Saber M 1989 *J. Mol. Evol.* **29** 513
[30] de Oliveira V M, Fontanari J F and Stadler P F 1999 *J. Phys. A: Math. Gen.* **32** 8793
[31] Derrida B 1997 *Physica D* **107** 186
[32] Bouchaud J-P and Mézard M 1997 *J. Phys. A: Math. Gen.* **30** 7997