

## FITNESS LANDSCAPES AND THE ANDREWS–CURTIS CONJECTURE

JERRY SWAN\*, GABRIELA OCHOA†  
and GRAHAM KENDALL‡

*Automated Scheduling, Optimisation  
and Planning (ASAP) Research Group  
School of Computer Science and Information Technology  
University of Nottingham, Jubilee Campus  
Wollaton Road, Nottingham NG8 1BB, UK*

\**jps@cs.nott.ac.uk*

†*gxo@cs.nott.ac.uk*

‡*gxk@cs.nott.ac.uk*

MARTIN EDJVET

*School of Mathematical Sciences  
University of Nottingham, University Park  
Nottingham, NG7 2RD, UK  
martin.edjvet@nottingham.ac.uk*

Received 2 December 2010

Accepted 5 August 2011

Published 23 February 2012

Communicated by Mark Sapir

Attempts have been made to eliminate some potential counterexamples to the Andrews–Curtis conjecture using the combinatorial optimization methods of blind-search and the genetic algorithms meta-heuristic. Breadth-first search with secondary storage is currently the most successful method, which raises questions regarding the inferior performance of heuristic search. In order to understand the underlying reasons we obtain fitness landscape metrics for a number of balanced presentations and draw conclusions regarding the likely effectiveness of other meta-heuristics.

*Keywords:* Andrews–Curtis conjecture; fitness landscapes; combinatorial optimization.

Mathematics Subject Classification 2010: 20-04

### 1. Introduction

The Andrews–Curtis conjecture was formulated in 1965 and is of interest in group theory and low-dimensional topology [2]. It states that “every balanced presentation of the trivial group can be transformed into the trivial presentation via some sequence of *AC-moves*”.

A finite presentation  $\langle g_1, \dots, g_m \mid r_1, \dots, r_n \rangle$  is said to be *balanced* if  $m$  is equal to  $n$ . The *trivial presentation* of the trivial group of rank  $r$  is the balanced presentation  $\langle g_1, \dots, g_r \mid g_1, \dots, g_r \rangle$ . The AC-moves are:

**AC1.**  $r_i \rightarrow r_i^{-1}$

**AC2.**  $r_i \rightarrow r_i r_j, i \neq j$

**AC3.**  $r_i \rightarrow g^{\mp 1} r_i g^{\pm 1}$ , for some generator  $g$

These are the regular Nielsen transformations together with conjugation of a relator by an arbitrary word in the generators via repeated application of AC3. There is also a variant of the conjecture featuring a “stabilization” move (i.e. introducing a new generator  $y$  along with the relator  $y = 1$ ) which is of particular interest to topologists. The AC-moves on the balanced presentations of rank  $n$  generate a group,  $AC_n$ , with corresponding action on the set of balanced presentations of rank  $n$ . In the absence of a proof of the conjecture, the status of counterexamples has attracted attention and there is interest in using computer search to eliminate them (families of potential counterexamples are given in [17, 19]). Given a potential counterexample  $P$ , the aim is to find a word  $w$  in  $AC_n$  that reduces  $P$  to the trivial presentation (or equivalently to find  $w^{-1}$  transforming the trivial presentation into  $P$ ). Computationally, this can be achieved by the discrete optimization technique known as state-space search, the most naïve exemplars of which are exhaustive enumeration and randomly generating candidate solutions. In [14], Miasnikov employed genetic algorithms to determine that all the well-known potential counterexamples with total length of relators at most 12 satisfy the conjecture. In particular he solved the  $m = 2$  case from the family of potential counterexamples  $\langle a, b \mid a^m = b^{m+1}, aba = bab \rangle$  due to Akbulut and Kirby [1]. In [15], Miasnikov and Myasnikov extend the length 12 results to include all balanced presentations with total length of relators at most 12. Havas and Ramsay subsequently employed breadth-first search modulo AC-equivalence and were able to further extend these results to the length 13 case [6].

For a balanced presentation of rank  $n$ , there are  $3n^2$  AC-moves, and hence  $(3n^2)^l$  move sequences of length  $l$ . Many move sequences will result in the same presentation so the value of  $(3n^2)^l$  is really only a rough upper bound. Even though the branching factor is thus somewhat constrained, it remains exponential and significant further progress employing main-memory alone is not possible. Bowman and McCaul [4] extended the breadth-first search approach by making extensive use of secondary storage indexed via an in-memory hashing scheme, solving all 10 presentations of the form  $\langle a, b \mid a = [a^m, b^n], b = [a^p, b^q] \rangle, m, n, p, q \in \mathbb{N}$  having relator length 14. By employing constraints on the maximum permissible total relator length for intermediate presentations, they also exhaustively enumerated the  $m = 3$  case for Akbulut–Kirby (currently the smallest potential counterexample) for total relator lengths from 10 to 17 inclusive, but were unable to find an Andrews–Curtis trivialization.

This paper seeks to better understand why breadth-first search is currently enjoying greater success than a heuristically informed technique such as genetic algorithms. We do this by analyzing the *fitness landscape* of some potential counterexamples. In the remainder of this paper, we apply a series of fitness landscape analysis techniques to highlight the reasons for the failure of heuristic search and reach some conclusions regarding the nature of those techniques that are likely to enjoy greater success than those applied to date.

## 2. Methodology

A ubiquitous definition of a fitness landscape [13] is a triple  $(S, f, d)$  where  $S$  is the set of states,  $f : S \rightarrow \mathbb{R}^+$  is the *fitness function* and  $(S, d)$  is a metric space. We elect to define landscapes via the triple  $(S, f, N)$  where  $N : S \rightarrow 2^S$  is a *neighborhood* of  $s$ , i.e. the set of all states directly reachable from  $s$ . The neighborhood induces a graph  $\Gamma$  with vertex set  $V = \{s \in S\}$  and edge set  $E = \{(s, t) \mid t \in N(s)\}$ . We then take  $d(s, s')$  to be the length of the shortest path between  $s$  and  $s'$ . In the following sections, we proceed to detail the components of the triple  $(S, f, N)$  used to define our landscapes and introduce some elementary fitness landscape metrics.

### 2.1. State-spaces and neighborhoods

Meta-heuristic state-spaces may be categorized as *perturbative* or *constructive*. In the former case the vertices of the state-space graph are complete solution representations, in the latter they are *partial* solutions, with the mapping from partial to complete solution being defined in a problem-specific manner. The breadth-first search approach of [6] is perturbative, while the genetic algorithms approach of [14] is constructive.

Let  $B_P$  be the graph generated from the action of all words in  $AC_n$  on some balanced presentation  $P$  of rank  $n$ . The vertices of this graph are balanced presentations of rank  $n$  and the edges are the AC-moves. If the Andrews–Curtis conjecture is true, when  $P$  is the trivial presentation of rank  $n$  the vertex set of this graph includes every balanced presentation of that rank. This is the space searched (modulo constraints on depth and relator length) by the breadth-first search implementations in [6] and [4]. Let  $N_B(s)$  be the corresponding neighborhood, i.e. the set of all presentations obtained from  $s$  via the action of a single AC-move.

Let  $AC_n^k$  be the set of all words in  $AC_n$  of length  $k$ . For dynamically variable  $k$ , this is the space searched by Miasnikov’s genetic algorithms (his mutation operators included insertion and deletion of elements in a word). For  $w \in AC_n^k$ , let  $H_1(w)$  be the neighborhood of all words in  $AC_n^k$  with Hamming distance 1 from  $w$ .

### 2.2. Fitness functions

The fitness function used by Miasnikov attempts to minimize the total length of relators. Havas and Campbell commented that this could not be expected to be a good heuristic measure since it behaves in a highly non-monotonic fashion along the

path to a solution. We compare this RELATOR-LENGTH fitness function with two additional fitness functions [5]: LONGEST-COMMON-SUBSTRING yields the length of the longest common substring between a candidate solution and the target presentation and SUBSTRING-EDIT-DISTANCE yields a measure of the number of substitutions, insertions and deletions required to transform the source presentation into the target. The fitness value for both these functions is determined to be the maximum (respectively minimum) over all pairwise assignments between candidate and target relators.

### 2.3. Fitness landscape metrics

Knowledge of the features of a search landscape is of interest both for predicting the performance and for improving the design of search algorithms. Statistical measures have been proposed for measuring fitness landscape properties: two of the most commonly used landscape analysis techniques (auto-correlation analysis and fitness–distance correlation) are described in detail below.

#### 2.3.1. Autocorrelation analysis

An important characteristic of a landscape is its *ruggedness*, which is related to the difficulty of an optimization problem for meta-heuristics. Weinberger [20] introduced a procedure to measure the ruggedness of a fitness landscape based on the *autocorrelation function* [8]. The idea is to generate a random walk of length  $T$ , on the landscape via neighboring points. At each step, the fitness of the solution encountered is recorded and a corresponding time series of fitness values is obtained. Thereafter, the autocorrelation function of the time series,  $\rho_i$  may be calculated. This theoretical autocorrelation function  $\rho_i$  can be empirically estimated by  $r_i$ :

$$r_i = \frac{\sum_{t=1}^{T-i} (f_t - \bar{f})(f_{t+i} - \bar{f})}{\sum_{t=1}^T (f_t - \bar{f})^2}, \quad (1)$$

where  $\bar{f}$  is the mean fitness of the  $T$  points visited, and  $i$  is the time lag or distance between points in the walk. The assumption underlying the use of this metric is that the landscape is *statistically isotropic*, i.e. that the correlation for a single random walk is representative of the landscape as a whole.

A related measure is the *correlation length* of a fitness landscape. Several authors have proposed approaches to measure this quantity [8, 11, 20]. Statistically the correlation length gives an indication of the largest “distance” (or time lag) between two points at which the value of one point can still provide information about the expected value of the other point [8]. In other words, the correlation length is the largest time lag  $i$  for which one can still expect some correlation between two points  $i$  steps apart. We use here the correlation length measure based on the estimated autocorrelation function [20]  $\ell = \frac{-1}{\ln(|r_1|)}$  for  $r_1 \neq 0$  (where  $r_1$  is defined

according to Eq. (1)). Correlation length is inversely proportional to ruggedness, and lower values are indicative of landscapes that are problematic for heuristic search.

### 2.3.2. Fitness–distance correlation

The most commonly used measure to estimate the global structure of fitness landscapes is the *fitness–distance correlation* (FDC) coefficient, proposed by Jones and Forrest [9]. It is used as a measure for problem difficulty in genetic algorithms. Given a set of points  $\{x_1, x_2, \dots, x_m\}$  and their fitness values, the FDC coefficient  $\varrho$  is defined as:

$$\varrho(f, d_{\text{opt}}) = \frac{\text{Cov}(f, d_{\text{opt}})}{\sigma(f)\sigma(d_{\text{opt}})}, \quad (2)$$

where  $\text{Cov}(\cdot, \cdot)$  denotes the covariance of two random variables and  $\sigma$  the standard deviation. The FDC determines how closely related are the fitness of a set of points and their distances to the nearest optimum in the search space (denoted by  $d_{\text{opt}}$ ). A value of  $\varrho = -1.0$  ( $\varrho = 1.0$ ) for maximization (minimization) problems indicates a perfect correlation between fitness and distance to the optimum, and thus predicts an easy search. On the other hand, a value of  $\varrho = 1.0$  ( $\varrho = -1.0$ ), means that with increasing fitness the distance to the optimum increases too, which indicates a deceptive and difficult problem. As suggested in [9], a value of  $\text{fdc} \leq -0.5$  ( $\text{fdc} \geq 0.5$ ) for maximization (minimization) problems is indicative of an easy problem.

A *fitness–distance plot* can also provide additional insight into the structure of the landscape [12]. This is done by plotting the fitness of points in the search space against their distance to an optimum solution. This type of analysis can be used to investigate not only the correlation between arbitrary points in the search space, but also the distribution of local optima within the search space. Should the global optimal be unknown, fitness–distance plots can work from a best-known solution, but this may yield a significantly different plot.

## 3. Analysis of Andrews–Curtis Landscapes

For our analysis, we considered the following families:

$$\text{HR}_m = \langle a, b \mid a^b = a^m, b^a = b^m \rangle \text{ as discussed in [6],} \quad (3)$$

$$\text{AK}_m = \langle a, b \mid a^m = b^{m+1}, aba = bab \rangle \text{ from [1].} \quad (4)$$

We determined values for the correlation lengths for both  $N_B(s)$  and  $H_1(w)$  neighborhoods under the action of the three fitness functions described above. We then present corresponding fitness–distance plots and fitness histograms for the RELATOR-LENGTH fitness function.

### 3.1. Autocorrelation analysis

Tables 1–3 give correlation lengths of the RELATOR-LENGTH fitness function on the  $N_B(s)$  neighborhood for  $HR_2$  and the  $m = \{2, 3\}$  cases for  $AK_m$ . The columns of Tables 1–3 are:

**Search-direction.** A sufficient condition for the success of forward searches is the reduction of  $n - 1$  relators to primitive elements [14]. We incorporate this condition for success into forward searches. For reverse searches we check modulo AC-equivalence to see whether the target presentation has been found.

**Whitehead?** The approach adopted in [14] was to supplement the AC-moves with *Whitehead automorphisms*. For generators  $g_i, g_j$ , with  $i \neq j$  they are given by:

$$g_i \rightarrow g_i^{-1}, \quad g_i \rightarrow g_i g_j^{\pm 1}, \quad g_i \rightarrow g_j^{-1} g_i g_j$$

Table 1. Correlation length of RELATOR-LENGTH for state-space  $B_P, P = HR_2$ .

Search-direction	Whitehead?	Total relator-length constraint	Correlation-length
Forward	No	15	13.8385
Forward	Yes	15	5.21098
Reverse	No	15	3.26104
Reverse	Yes	15	5.06388
Forward	No	19	3.09834
Forward	Yes	19	6.97969
Reverse	No	19	3.29097
Reverse	Yes	19	4.92711

Table 2. Correlation length of RELATOR-LENGTH for state-space  $B_P, P = AK_2$ .

Search-direction	Whitehead?	Total relator-length constraint	Correlation-length
Forward	No	17	2.26545
Forward	Yes	17	6.21152
Reverse	No	17	3.52197
Reverse	Yes	17	4.30575
Forward	No	21	3.00571
Forward	Yes	21	5.41039
Reverse	No	21	2.65021
Reverse	Yes	21	5.09715

Table 3. Correlation length of RELATOR-LENGTH for state-space  $B_P, P = AK_3$ .

Search-direction	Whitehead?	Total relator-length constraint	Correlation-length
Forward	No	20	2.24334
Forward	Yes	20	2.65731
Reverse	No	20	3.50911
Reverse	Yes	20	4.78769
Forward	No	24	3.13675
Forward	Yes	24	2.75453
Reverse	No	24	3.12792
Reverse	Yes	24	4.85375

The addition of Whitehead automorphisms allows for shorter proofs, although they also increase the branching factor of the search tree.

**Total relator-length constraint.** As observed in [6], the addition of constraints on the inclusion of a state in the breadth-first search tree makes the state-space finite (although the search is then no longer guaranteed to find a solution if one exists).

It might be expected that the additional success criterion for forward searches would reduce the predictive quality of their associated fitness function due to the discontinuity it imposes on the fitness gradient (the “needle in a haystack” effect). In fact, it was determined that the opposite is true for all of the random-walk correlations we performed.

Correlation length can be seen to decrease if the pruning constraint of total relator-length is relaxed. Correlation lengths were also determined for LONGEST-COMMON-SUBSTRING and SUBSTRING-EDIT-DISTANCE but were found in all cases to give values close to that of RELATOR-LENGTH. Taken together, these observations imply that none of these fitness functions afford a useful objective measure for searching the unconstrained landscape.

Tables 4 and 5 give correlation lengths of the RELATOR-LENGTH fitness function for the  $H_1(w)$  neighborhood for  $w \in \text{AC}_2^k$  as applied to  $\text{HR}_2$  and  $\text{AK}_2$ . The choice of  $k$  is motivated by the length of the shortest path to a solution (determined by breadth-first search as described in more detail in Sec. 3.2). The constraint on total relator-length is omitted for the  $\text{AC}_n^k$  state-space, since for low values of relator-length, most solutions in this state-space will fail to meet this constraint. Once again, no significant difference was observed in the values for the other two fitness functions.

Table 4. Correlation length of RELATOR-LENGTH in neighborhood  $H_1$  of  $\text{AC}_2^5$  for  $\text{HR}_2$ .

Search-direction	Whitehead?	Correlation-length
Forward	No	3.25263
Forward	Yes	3.66581
Reverse	No	1.7498
Reverse	Yes	1.65063

Table 5. Correlation length of RELATOR-LENGTH in neighborhood  $H_1$  of  $\text{AC}_2^{10}$  for  $\text{AK}_2$ .

Search-direction	Whitehead?	Correlation-length
Forward	No	5.75898
Forward	Yes	4.29469
Reverse	No	3.55046
Reverse	Yes	3.10686

With the exception of the outlying value of 13.8385, all these landscapes are moderately rugged [18]. As observed in [10], there is common consensus that the number of local optima increase with ruggedness and we proceed to investigate this in the next section.

### 3.2. *Fitness–distance correlation*

As discussed above, the main purpose of fitness–distance analysis is to explore the distribution of local optima in the state-space. We examined HR<sub>2</sub> and AK<sub>2</sub> instances using forward search, RELATOR-LENGTH fitness function and without Whitehead automorphisms. As discussed above, it is desirable to apply the FDC to an optimal solution. We therefore employed breadth-first search to find AC trivializations for HR<sub>2</sub> and AK<sub>2</sub> using the canonical form described in [4]. The searches were performed on a 3 GHz Pentium® IV PC with 2 GB RAM. The proof sequence for HR<sub>2</sub> of  $(BabAA, AbaBB) \rightarrow (BabAA, baBBA) \rightarrow (BabAA, aBBAb) \rightarrow (BabAA, BA) \rightarrow (BabAA, ab) \rightarrow (BabAA, b)$  of length 5 was obtained in 4 seconds. The proof sequence for AK<sub>2</sub> of  $(aaBBB, abaBAB) \rightarrow (aaBBB, babABA) \rightarrow (aaBBB, abABaBB) \rightarrow (bbbAA, abABaBB) \rightarrow (bAbABa, abABaBB) \rightarrow (bAbABa, bABaBBa) \rightarrow (AbaBaB, bABaBBa) \rightarrow (AbaBaB, BaBaB) \rightarrow (AbaBaB, bAbAb) \rightarrow (AbaBaB, bAb) \rightarrow (AbaBaB, b)$  of length 10 was obtained in 23 seconds.

The states of the FDC are thus sequences of AC-moves in AC<sub>2</sub><sup>k</sup>, with *k* being given by the respective lengths of the above proof sequences. Since the rank is equal to 2 in both cases, the number of AC-moves is equal to 12 and the state-spaces have sizes 12<sup>5</sup> and 12<sup>10</sup> for the HR and AK instances, respectively. Hamming distance was again used as the distance metric between sequences. For each instance, 5000 random candidate solutions were generated. For each of these, 5000 local optima were produced by running a next-descent local search algorithm (hill-climber). This algorithm (detailed in Fig. 1) accepts solutions of improving fitness, and uses a 1-opt neighborhood, which uniformly at random selects a position in the sequence

<pre> <b>Procedure</b> Next-Descent(<i>P</i> : Balanced-Presentation, <i>s</i> : sequence of AC-moves)   <math>\pi \leftarrow</math> permutation of indices of <i>s</i>   <b>for</b> <math>i \leftarrow 1</math> to <i>Length</i>(<i>s</i>) <b>do</b>     <math>s' \leftarrow</math> mutate AC-move <math>\pi[i]</math> of <i>s</i>     <math>P' \leftarrow</math> apply <math>s'</math> to <i>P</i>     <b>if</b> RELATOR-LENGTH(<math>P'</math>) &lt; RELATOR-LENGTH(<i>P</i>) <b>then</b>       <math>s \leftarrow s'</math>     <b>end if</b>   <b>end for</b> </pre>
---

Fig. 1. Pseudo-code for the algorithm iterated to a local optimum. The parameter *P* is the source presentation and *s* represents the incumbent solution.



Table 6. Results of the fitness–distance analysis.

Instance	$k$	Opt	min $d_{\text{opt}}$	max $d_{\text{opt}}$	avg $d_{\text{opt}}$	avg $d_{\text{loc}}$	$N$	fdc
HR <sub>2</sub>	5	3	0	5	3.34	4.16	2097	0.61
AK <sub>2</sub>	10	4	0	10	7.12	8.54	4992	−0.02

and replaces it with a randomly selected AC-move. The results of the analysis are summarized in Table 6. The first column displays the name of the instance, the second the sequence length ( $k$ ), and the third the fitness of the best-known solution (Opt). Columns four to six display the minimum, maximum and average distance of the local optima to the global optimum (min  $d_{\text{opt}}$ , max  $d_{\text{opt}}$  and avg  $d_{\text{opt}}$ ). Columns seven to nine show the average distance between local optima (avg  $d_{\text{loc}}$ ), the number of distinct local optima out of 5000 ( $N$ ), and the fitness–distance correlation coefficient.

For both instances, the local optima are distributed across the whole search space. As the maximum and average distance to the optimum indicate, they are not clustered around the global optimum. The maximum distance in both cases is the largest possible ( $k$ , the sequence length), and the average distance is greater than  $k/2$ . The average distance between local optima is larger than the average distance between local optima and the global optimum, suggesting that the global optimum is centrally located among the optima. The HR<sub>2</sub> case shows a positive and large fitness–distance correlation, whereas AK<sub>2</sub> shows a small negative correlation. The number of distinct local optima is less than the half of the sample size (5000) for HR<sub>2</sub>, whereas for AK<sub>2</sub> most local optima are distinct. The fitness–distance plots for the two instances are shown in Fig. 2. For all the local optima, the distance to the global optimum is plotted against the fitness. In both instances, the local optima are distributed widely across the  $k$  possible distances to the global optimum. Moreover, for some fitness values, there are several local optima at different distances, while




Fig. 2. Fitness–distance plots: (a) HR<sub>2</sub>, (b) AK<sub>2</sub>.

for most fitness values there is not even a single one, which leads to large gaps in fitness of the local optima. For example, most local optima in  $AK_2$  have fitness values in the range 10 to 15. In both cases, this hints at the existence of large plateaus in the search landscape.

In order to illustrate the distribution of fitness values of both randomly generated candidate solutions and of local optima, Figs. 3 and 4 show histograms of these values on  $HR_2$  and  $AK_2$  respectively. Notice that the range of fitness values naturally reduces for the local optima as compared to the randomly generated solutions and this reduction is more noticeable in  $AK_2$  (Fig. 4). In both instances, the most remarkable feature of these plots is that the majority of local optima share a small fraction of the fitness function values, confirming the existence of large plateaus away from the global optimum. This is again more noticeable in  $AK_2$  (see Fig. 4(b)), where most local optima have a fitness value of 11. This is a clear




Fig. 3. Fitness values histogram for  $HR_2$ . (a) Randomly generated points, (b) local optima.




Fig. 4. Fitness values histogram for the  $AK_2$ . (a) Randomly generated points, (b) local optima.

indication of the difficulty of searching on these landscapes since such plateaus can act as traps for any stochastic local search heuristic.

### 3.3. Landscape neutrality

The fitness–distance plots of the previous section suggest that a significant obstacle to search progress is the presence of extensive plateaus in the fitness landscape. Formally, we can define a plateau as the transitive closure of neighboring vertices of equal fitness. One means of depicting these neutral features of a landscape is via the use of *plateau-connection graphs* [7]. A plateau-connection graph may be obtained either from exhaustive enumeration or from sampling the landscape. A plateau is said to be *open* if and only if it contains some vertex having a neighbor with a better fitness value, otherwise it is said to be *closed*.

Figure 5 shows those connected components of the plateau-connection graph for  $B_P, P = \text{HR}_2$  that lead to a solution. Open plateaus are denoted by an ellipse, closed plateaus by a rectangle and solutions are denoted by triangles. Nodes are annotated as “# plateau-size @ fitness”. Figure 6 does the same for the  $H_1, k = 5$  case of  $\text{HR}_2$ . Figure 6 is illustrative in that it contains a closed plateau that can act as a *sink*, i.e. the search will remain within its basin of attraction unless a succession of unimproving moves are accepted.

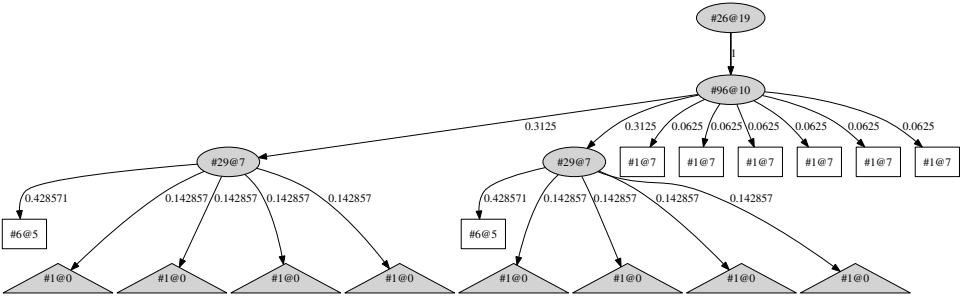


Fig. 5. Connected components of the plateau-connection graph for  $B_P, P = \text{HR}_2$ .

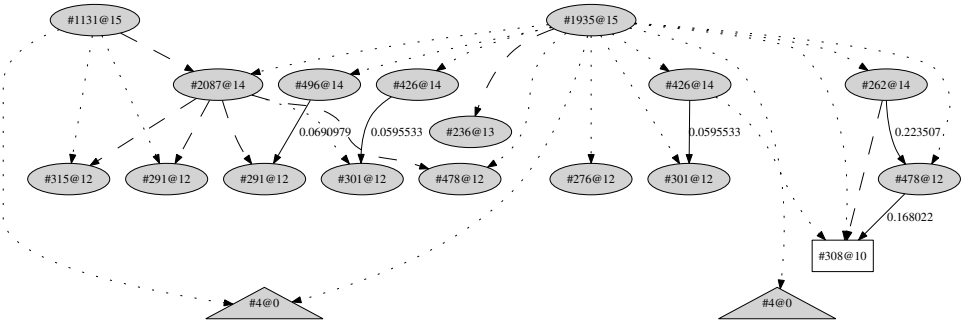


Fig. 6. Connected components of the plateau-connection graph for the  $H_1, k = 5$ , case of  $\text{HR}_2$ .

## 4. Conclusion

We measured local and global properties from the landscapes obtained for both perturbative and constructive state-space representations ( $B_P$  and  $AC_n^k$  respectively). We investigated the effect of some alternative fitness functions and concluded that none of them provide significant information about the larger-scale structure of the landscape.

By plotting fitness histograms and plateau-connection graphs, we were further led to conclude that the existence of large plateaus in the search space is the most significant contribution to the intractability of potential counterexamples. The inherent symmetry of the search-space implies the existence of much deeper basins of this type for larger problem instances. This is the likely explanation for Miasnikov’s genetic algorithm requiring an uncommonly high (95%) mutation probability, effectively degenerating into random search.

This suggests that the standard palette of meta-heuristics might also perform no better than random search unless augmented by some mechanism for “reacting on the fitness function” e.g. by imposing a gradient on plateaus via some tie-breaking function of the internal structure of a presentation [3]; via the automated generation of fitness functions that yield more favorable landscapes, or via an adaptive penalty mechanism (as used in *guided local search* [16]).

## References

- [1] S. Akbulut and R. Kirby, A potential smooth counterexample in dimension 4 to the Poincare conjecture, the Schoenflies conjecture, and the Andrews–Curtis conjecture, *Topology* **24** (1985) 375–390.
- [2] J. J. Andrews and M. L. Curtis, Free groups and handlebodies, *Proc. Amer. Math. Soc.* **16**(2) (1965) 192–195.
- [3] R. Battiti, M. Brunato and F. Mascia, *Reactive Search and Intelligent Optimization*, Operations Research/Computer Science Interfaces Series, Vol. 45 (Springer-Verlag, 2008).
- [4] R. S. Bowman and S. B. McCaul, Fast searching for Andrews–Curtis trivializations, *Experiment. Math.* **15**(3) (2006).
- [5] D. Gusfield, *Algorithms on Strings, Trees, and Sequences: Computer Science and Computational Biology* (Cambridge University Press, 2007).
- [6] G. Havas and C. Ramsay, Breadth-first search and the Andrews–Curtis conjecture, *Internat. J. Algebra Comput.* **13**(1) (2003) 61–68.
- [7] H. Hoos and T. Sttzele, *Stochastic Local Search: Foundations & Applications* (Morgan Kaufmann San Francisco, CA, USA, 2004).
- [8] W. Hordijk, A measure of landscapes, *Evol. Comput.* **4**(4) (1997) 335–360.
- [9] T. Jones, Crossover, macromutation, and population-based search, in *Proc. of the Sixth Int. Conf. Genetic Algorithms*, ed. L. J. Eshelman (Morgan Kaufmann, San Francisco, CA, 1995), pp. 73–80.
- [10] L. Kallel, B. Naudts and C. R. Reeves, *Properties of Fitness Functions and Search Landscapes* (Springer-Verlag, London, 2001), pp. 175–206.
- [11] B. Manderick, M. de Weger and P. Spiessens, The genetic algorithm and the structure of the fitness landscape, in *Proc. of the 4th Int. Conf. Genetic Algorithms* (Morgan Kaufmann, 1991), pp. 143–150.

- [12] P. Merz and B. Freisleben, Fitness landscapes, memetic algorithms, and greedy operators for graph bipartitioning, *Evol. Comput.* **8**(1) (2000) 61–91.
- [13] P. Merz and B. Freisleben, Fitness landscapes and memetic algorithm design, in *New Ideas in Optimization* (McGraw-Hill, 1999), pp. 245–260.
- [14] A. D. Miasnikov, Genetic algorithms and the Andrews–Curtis conjecture, *Internat. J. Algebra Comput.* **9** (1999) 671–686.
- [15] A. D. Miasnikov and A. G. Myasnikov, Balanced presentations of the trivial group on two generators and the Andrews–Curtis conjecture, in *Groups and Computation III*, Vol. 23, eds. W. Kantor and A. Seress (de Gruyter Berlin, 2003), pp. 257–263.
- [16] P. Mills and E. P. K. Tsang, Guided local search for solving SAT and weighted MAX-SAT problems, *J. Automat. Reason.* **24**(1/2) (2000) 205–223.
- [17] A. G. Myasnikov, A. G. Myasnikov and V. Shpilrain, On the Andrews–Curtis equivalence, *Contemp. Math.* **296** (2002) 183–198.
- [18] W. Rowe, M. Platt, D. C. Wedge, P. J. Day, D. B. Kell and J. Knowles, Analysis of a complete DNA-protein affinity landscape, *J. R. Soc. Interface* **7**(44) (2010) 397–408.
- [19] P. Schupp and C. F. Miller III, Some presentations of the trivial group, in *Groups, Languages and Automata*, ed. R. Gilman, Contemporary Mathematics, Vol. 250 (American Mathematical Society, Providence, RI, 1999), pp. 113–115.
- [20] E. D. Weinberger, Correlated and uncorrelated fitness landscapes and how to tell the difference, *Biol. Cybernetics* **63** (1990) 325–336.