

Understanding Phase Transitions with Local Optima Networks: Number Partitioning as a Case Study

Gabriela Ochoa¹(✉), Nadarajen Veerapen¹, Fabio Daolio¹, and Marco Tomassini²

¹ Computing Science and Mathematics, University of Stirling, Stirling, Scotland, UK
gabriela.ochoa@cs.stir.ac.uk

² Faculty of Business and Economics, Information Systems Department, University of Lausanne, Lausanne, Switzerland

Abstract. Phase transitions play an important role in understanding search difficulty in combinatorial optimisation. However, previous attempts have not revealed a clear link between fitness landscape properties and the phase transition. We explore whether the global landscape structure of the number partitioning problem changes with the phase transition. Using the local optima network model, we analyse a number of instances before, during, and after the phase transition. We compute relevant network and neutrality metrics; and importantly, identify and visualise the funnel structure with an approach (*monotonic sequences*) inspired by theoretical chemistry. While most metrics remain oblivious to the phase transition, our results reveal that the funnel structure clearly changes. Easy instances feature a single or a small number of dominant funnels leading to global optima; hard instances have a large number of suboptimal funnels attracting the search. Our study brings new insights and tools to the study of phase transitions in combinatorial optimisation.

1 Introduction

It has been recognised that phase transitions play an important role in analysing combinatorial optimisation problems; yet a clear link between fitness landscape structure and the phase transition phenomenon is still lacking. We use the local optima networks model to analyse and visualise the global structure of Number Partitioning fitness landscapes.

The Number Partitioning Problem (NPP) is defined as follows. Given a set of N positive numbers $L = \{r_1, r_2, \dots, r_N\}$, find a partition $A \cup B = L$ such that the partition difference

$$D = \left| \sum_{r_i \in A} r_i - \sum_{r_i \in B} r_i \right|$$

is minimised. The decision version of the NPP belongs to the class of NP-complete problems which appear to require a super-polynomial amount of computation time in the instance input size [1,2]. NP-hard optimisation problems are

at least as hard as the corresponding decision problems. Many important practical optimisation problems are NP-hard and it is thus important to find efficient approximate methods to solve them. Though requiring exponential time to be solved in the worst case, several hard constraint-satisfaction problems show an instance-dependent *computational phase transition*, meaning that below some critical point, instances are typically easy to solve while they become hard to solve above such a point. Well-known examples are the Boolean Satisfaction Problem (SAT) [1,3], the Graph-Colouring problem [1,4], and the NPP [1,5]. The control parameter is a problem-dependent quantity that must be suitably defined; for example, in SAT, it is the ratio of the number of clauses to the number of variables, and the phase transition phenomenon has been observed in both random and structured instances [6,7].

For problems undergoing a computational hardness phase transition it is of interest to understand how it arises and what are the problems' features that characterise the transition. The original methodology was developed by physicists and it is based on the statistical mechanics approach to physical phase transitions such as the ferromagnetic/paramagnetic transition. An introduction to this rather technical field as applied to hardness phase transitions can be found in [8]. We consider, instead, fitness landscape analysis as a tool for revisiting the phase transition phenomenon. In particular, the Local Optima Network (LON) model [9,10]. Local optima networks compress the whole search space into a graph, where nodes are local optima and edges are transitions among them with a given search operator. Local optima are key features of fitness landscapes as they can be seen as obstacles for reaching high quality solutions. The local optima networks model emphasises the number, distribution and most importantly, the connectivity pattern of local optima in the underlying search space. They are therefore an ideal tool for modelling and visualising the global structure of fitness landscapes. Among local optima network metrics, we particularly study the presence and distribution of so-called funnels in the landscape.

The term 'funnel' was introduced in the protein folding community to describe "a region of configuration space that can be described in terms of a set of downhill pathways that converge on a single low-energy structure or a set of closely-related low-energy structures" [11]. It has been suggested that the energy landscape of proteins is characterised by a single deep funnel, a feature that underpins their ability to fold to their native state. In contrast, some shorter polymer chains (polypeptides) that misfold are expected to have other funnels that can act as traps. Energy landscapes are conceptually related to fitness landscapes, and funnel structures have also been studied in heuristic continuous optimisation [12,13], and more recently in combinatorial optimisation [14,15,16].

The next section overviews previous and related work. Section 3 presents relevant definitions and algorithms related to local optima networks. Section 4 presents our fitness landscape analysis and visualisation. Finally, Section 5 summarises our main findings and suggests directions of future work.

2 Background and Related Work

2.1 The number partitioning fitness landscape

The existence of the NPP hardness phase transition was first demonstrated numerically by Gent and Walsh [5], who introduced the control parameter k and estimated the transition point to occur around $k_c = 0.96$. The control parameter k corresponds to the number of significant bits in the encoding of the input numbers r_i divided by N (the instance size), specifically $k = \log_2(M)/N$, where M is the largest number in the set $L = \{r_1, r_2, \dots, r_N\}$. For $\log_2(M)$ and N tending to infinity, the transition occurs at the critical value of $k_c = 1$, such that for $k < 1$, there are many perfect partitions with probability tending to 1, whereas for $k > 1$, the number of perfect partitions drops to zero with probability tending to 1 [5].

Further studies within the physics community, have confirmed the existence of the NPP phase transition characterising it rigorously [17,18]. However, they provide no direct answer to the question of what features of the corresponding fitness landscapes, if any, are responsible for the widely different observed behaviour. A step in this direction was taken by Fontanari et al. [19] who studied various landscape features before and after the transition, in particular considering *barrier trees* [20]. However, they were unable to find any effect of the transition on barrier tree features and other landscape metrics. Likewise, Alyahya and Rowe [21] performed an exhaustive statistical analysis of NPP landscapes for instances of size $N = 20$ and several number distributions, but did not find any significant correlation between most landscape features and easy or hard instances. They observed differences only in the number of global optima, which is high before the phase transition and low after it, and in the existence of neutral networks which are abundant in the easy phase and tend to disappear in the hard phase.

Considering these studies and given the lack of a clear picture, we decided to investigate additional landscape features based on local optima networks, and the recently proposed approach to identifying multiple funnels in combinatorial search spaces.

2.2 Multiple funnels in combinatorial landscapes

The big-valley hypothesis [22] suggests that on the travelling salesman problem (TSP) and other combinatorial optimisation problems, local optima are not randomly distributed, instead they are clustered around one central global optimum. Recent studies on TSP landscapes, however, have revealed a more complex picture [23,15,14]. The big-valley seems to decompose into several sub-valleys or multiple funnels. This helps to explain why certain iterated local search heuristics can quickly find high-quality solutions, but fail to consistently find the global optimum.

The procedure for identifying funnels on the TSP has evolved in recent work, ranging from visual inspection of the fitness distance correlation plots [23], con-

nected components in the local optima networks [14], and 3D LON visualisation [15]. We propose here to use the notion of *monotonic sequences* from theoretical chemistry [24], which describes a sequence of local minima where the energy of minima is always decreasing. We adapt this notion to the context of fitness landscapes and consider a monotonic sequence as a sequence of local optima where the fitness (costs) of solutions is non-deteriorating. The set of monotonic sequences leading to a particular optimum has been termed ‘basin’ [24], ‘monotonic sequence basin’ [25] and ‘super-basins’ [26], in the theoretical chemistry literature. We chose here to call them ‘funnel basins’ or simply ‘funnels’ borrowing from the protein folding literature. We can distinguish the *primary* funnel, as the one involving monotonic sequences that terminate at the global optimum (there can be more than one). The primary funnel is separated from other neighbouring secondary funnels by transition states laying on a so-called ‘primary divide’ [24]. Above such a divide, it is possible for a local optima to belong to more than one funnel through different monotonic sequences.

The presence of multiple funnels has also been recently observed on binary search spaces (*NK* landscapes) [16], where the authors observed a connection between groupings (communities) in local optima networks and the notion of funnels. Results confirm that landscapes consists of several clusters and the number of clusters increases with the epistasis level. A higher number of clusters leads to a higher search difficulty, measured by the empirical success rate of an iterated local search implementation. The success rate was also found to strongly correlate with the size of the cluster containing the global optimum.

There is evidence of clustering of solutions in Random Satisfiability problems [27] but a study of the funnel structure of SAT instances has not yet been conducted.

3 Definitions and Algorithms

This section overviews the definitions and algorithms constituting the local optima network model for the number partitioning problem.

3.1 Preliminaries

Fitness Landscape. A landscape [28] is a triplet (S, V, f) where S is a set of potential solutions, i.e., a search space; $V : S \rightarrow 2^{|S|}$, a neighbourhood structure, is a function that assigns to every $s \in S$ a set of neighbours $V(s)$, and $f : S \rightarrow \mathbb{R}$ is a fitness function that can be pictured as the *height* of the corresponding solutions.

In our study, the search space is composed of binary strings of length N , therefore its size is 2^N . The neighbourhood is defined as the 1-move or bit-flip operation, but definitions can be generalised to larger neighbourhoods.

Neutral neighbour. A neutral neighbour of s is a neighbour configuration x with the same fitness $f(s)$.

$$V_n(s) = \{x \in V(s) \mid f(x) = f(s)\}$$

The neutral degree of a solution is the number of its neutral neighbours. A fitness landscape is neutral if there are many solutions with high neutral degree. The landscape is then composed of several sub-graphs of configurations with the same fitness value.

Plateau. A plateau, also known in the literature as a neutral network [29,30], is a set of connected configurations with the same fitness value. Two vertices in a plateau are connected if they are neutral neighbours, that is, if they differ by one point mutation. With the bit-flip mutation operator, for all solutions x and y , if $x \in V(y)$ then $y \in V(x)$. So in this case, the plateaus are the equivalence classes of the relation $R(x, y)$ iff $(x \in V(y) \text{ and } f(x) = f(y))$.

Local optimum. A local optimum, which in the NPP case is a minimum, is a solution s^* such that $\forall s \in V(s^*), f(s^*) \leq f(s)$. Notice that the inequality is not strict, in order to allow the treatment of the neutral landscape case.

In the presence of neutrality, local minima are identified by a stochastic hill-climber h that, starting from any solution s , chooses the next best-improving mutant at each iteration by splitting ties at random, until convergence on a local optimum plateau.

Local optimum plateau. A plateau is a local optimum if all its configurations are local optima.

3.2 Local Optima Networks

In order to construct the networks, we need to define their nodes and edges. Nodes are local optima and edges represent escape probabilities. Local optima networks for neutral landscapes have been studied before by Verel et al. [10]; we borrow their notation and definitions, but name a sequence of connected solutions with the same fitness as *plateaus* rather than as *neutral networks*, to avoid confusion with the local optima network terminology.

Since we are interested in determining the landscape’s funnel structure using the notion of monotonic sequences, we only consider transitions between local optima where fitness is non-deteriorating. This leads to a variant of the model which we term *Monotonic Local Optima Networks (M-LON)*. Furthermore, our experiments revealed that neutrality is also present at the level of local optima transitions, that is, there are connected components in the M-LON which share the same fitness value. This leads us to define an even coarser model of the landscape, where these M-LON plateaus are compressed into single nodes, we termed this new model *Compressed Monotonic Local Optima Networks (CM-LON)*. Relevant formal definitions are given below.

LON nodes. The set of local optimum plateaus (formed of one or more local optima), $LOp = \{lop_1, lop_2, \dots, lop_n\}$ corresponds to the node set of the local optima network. The *basin of attraction* of a lop_i is the set of solutions $b_i = \{s \in S \mid h(s) = lop_i \text{ with probability } p_i(s) > 0\}$ and its size is $|b_i| = \sum_{s \in S} p_i(s)$.

Monotonic Edges. The set of monotonic edges, ME is defined according to a distance function d (minimal number of moves between two solutions), and a positive integer $D > 0$. Edges account for the chances of jumping from a local optimum plateau lop_i into the basin of a non-deteriorating local optimum plateau lop_j after a controlled perturbation. Namely, if we perturb a solution $s \in lop_i$ by applying D random moves, we obtain a solution s' that will belong to another basin b_j with probability p_j : that is, $h(s') = lop_j$ with probability p_j . The probability to go from s to b_j is then $p(s \rightarrow b_j) = \sum_{s' \in b_j} p(s \rightarrow s')p_j(s')$, where $p(s \rightarrow s') = P(s' \in \{z \mid d(z, s) \leq D\})$ is the probability for s' to be within D moves from s and can be evaluated in terms of relative frequency. Therefore, we can draw an edge e_{ij} between lop_i and lop_j with weight $w_{ij} = p(lop_i \rightarrow b_j) = \frac{1}{|lop_i|} \sum_{s \in lop_i} p_i(s)p(s \rightarrow b_j)$.

Monotonic Local Optima Network (M-LON). The weighted, oriented local optima network $M-LON=(LOp, ME)$ is the graph where the nodes $lop_i \in LOp$ are the local optimum plateaus, and there is an edge $e_{ij} \in ME$, with weight w_{ij} , between two nodes lop_i and lop_j if $w_{ij} > 0$.

M-LON plateau. Is a set of connected nodes in the M-LON with the same fitness value. Two nodes are connected if there is a monotonic edge between them.

Compressed LON nodes. The set of M-LON plateaus, $CLOp = \{clop_1, clop_2, \dots, clop_n\}$ corresponds to the node set of the compressed local optima network.

Compressed Monotonic Local Optima Network (CM-LON). The weighted, oriented local optima network $CM-LON = (CLOp, ME)$ is the graph where the nodes $clop_i \in CLOp$ are the M-LON plateaus. Weighted edges correspond to the aggregation of the multiple edges from nodes in a plateau to single edges in the compressed network. The weights of the multiple edges are added to constitute the weight of the mapped edge.

3.3 Detecting the funnel structures

To detect the funnel structures we first identify the funnels' 'ends' or 'bottoms'. To do so, we take advantage of the Compressed Monotonic Local Optima Networks. CM-LONs are directed graphs without loops. In a directed graph, one can distinguish the outdegree (number of outgoing edges) from the indegree (number of incoming edges); a *source* node is a node with indegree zero, while a *sink* node is a node with outdegree zero. We consider the CM-LONs sinks as the funnel bottoms.

We thus define the *funnel sinks* as the CM-LON nodes without outgoing edges. Once the funnel sinks are detected, we can proceed to identify the funnel basins (see Algorithm 1). This is done by finding all nodes in the CM-LON graph which are reachable from each funnel sink. Breadth-First-Search is used for this purpose. The set of unique nodes in the combined paths to a given funnel

sink corresponds to the funnel basin. The cardinality of this set corresponds to the funnel size. Notice that the membership of a node to a funnel might be overlapping, that is, a node may belong to more than one funnel, in that there are paths from that node to more than one funnel sink. The relative size of the primary funnel (or any other secondary funnel) is calculated as its size divided by the total number nodes in the graph.

Data: CM-LON: Compressed monotonic local optima network, S : funnel sinks

Result: $bsizes$: funnel basin sizes vector, $basins$: funnel basins vector

```

i ← 0
for  $s \in S$  do
     $basin[i] \leftarrow \text{breadthFirstSearch}(\text{CM-LON}, s)$ 
     $bsize[i] \leftarrow \text{length}(fbasin[i])$ 
     $i \leftarrow i + 1$ 
end

```

Algorithm 1: Identifying funnel basins.

4 Results and Analysis

One advantage of modelling landscapes as complex networks is the possibility of visualising them. After describing the experimental setting, we visualise a set of selected instances before, during, and after the phase transition. We continue with a study of local optima network metrics, including the new set of funnel measurements, and explore how they relate to the phase transition.

4.1 Experimental setting

In order to minimise the influence of the random creation of landscapes, we considered 30 different and independent landscapes for each parameter combination: N and k . Measurements consider the distribution of values across these 30 landscapes. The empirical study considers $N \in \{10, 15, 20\}$, where $N = 20$ is the largest possible value allowing practical exhaustive enumeration of the search space. The parameter k was varied from 0.4 to 1.2 in steps of 0.1. For each landscape, we extract the full local optima network using code adapted from Daolio et al. [31,32]. We then construct both the monotonic local optima networks (M-LON) and the compressed monotonic local optima networks (CM-LON). When extracting the local optima networks, we set the parameter D for the maximum escape distance to $D = 2$.

4.2 Visualisation

We visualise CM-LONs for selected instances with $N = 15$ and $k \in \{0.4, 0.6, 0.8, 1.0\}$. Due to space constraints, the instance with $k = 1.2$ is not shown, but it reflects a similar structure to that of $k = 1.0$

Network plots were produced using the R statistical language together with the `igraph` and `rgl` packages. Graph layouts consider *force-directed* methods. Networks are decorated to reflect features relevant to search dynamic. Red nodes correspond to global sinks, while blue nodes to suboptimal sinks; all other nodes are grey. An edge’s width is proportional to its weight, which indicates the probability of transitions. That is, the most probable transitions are thicker in the plots.

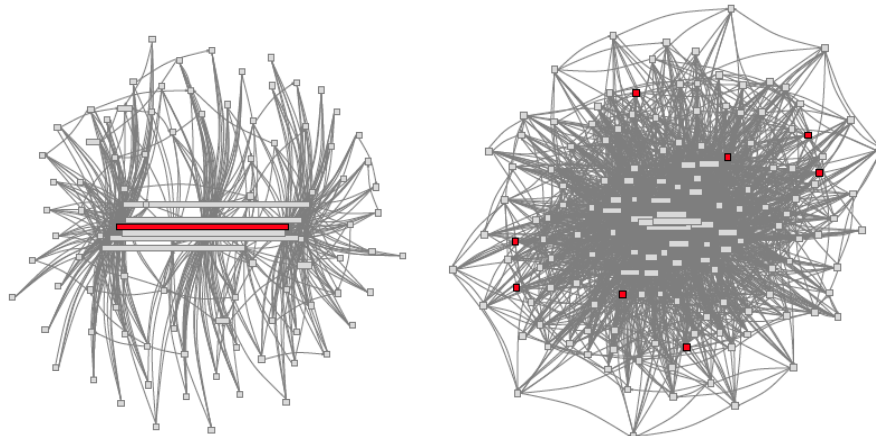
We explored two ways of visualising nodes. First, as rectangles (Figure 1) with lengths proportional to plateau sizes (i.e. the number of single local optima within a plateau). As the plots in Fig. 1 illustrate, for low values of k the landscape global optima form a large plateau, and there are several other large plateaus in the vicinity. With increasing k , the plateaus shrink, with nodes becoming single local optima for $k \geq 8$. Neutrality at the optima network level is, therefore, high for low values of k , gradually decreases with intermediate values of k and finally disappears for $k \geq 0.8$.

A second alternative is to visualise nodes with sizes proportional to their incoming strength (weighted incoming degree), as in Figure 2. Incoming strength is relevant as it reflects the extent to which a node ‘attracts’ the search dynamics; that is, it conveys the combined probability of a stochastic search process reaching it. We present both 2D and 3D images. In the 3D visualisations, the x and y coordinates are determined by the force-directed graph layout algorithm; while fitness is visualised as the z coordinate. This provides a clearer representation of the funnel and sink concepts, bringing an almost tangible aspect to these metaphors.

As Figure 2 illustrates, for $k = 0.4$ there is a single funnel structure easily guiding the search to the single global optimum. For $k = 0.6$, a single dominant central structure is still visible, but several different unconnected global optima now stem out from it. For $k \leq 0.6$ only optimal (red) sinks are observed, indicating that instances are easy to solve. When k increases over 0.6, suboptimal (blue) sinks start to emerge; initially only a few of them, but the number increases with increasing k . The number of optimal (red) sinks decreases and rapidly becomes only two. Search thus become harder, as can be inferred from the 2D and 3D visualisations of the landscape with $k = 1.0$; 16 blue sinks are observed and their combined incoming strength exceeds that of the 2 red sinks. Moreover, as indicated by the 3D image, some suboptimal blue sinks are deep, that is, they are close in fitness to the optimal solution.


4.3 Metrics

Due to space constraints, we can only visualise a few examples. Therefore, we turn to the statistical analysis of the complete dataset. Many features can be




(a) $N = 15, k = 0.4$

(b) $N = 15, k = 0.6$




(c) $N = 15, k = 0.8$





(d) $N = 15, k = 1.0$



Fig. 1: Local optima networks (CM-LONs) for selected NPP instances with $N = 15$. Nodes are local optimum plateaus visualised as rectangles with length proportional to their size (i.e. number of single local optima on them). Long rectangles indicate large plateaus, while squares indicate single local optima. For $k = 0.4$ the whole network is visualised, while for $k \in \{0.6, 0.8, 1.0\}$, the fittest part of the network is shown.



(a) $N = 15, k = 0.4$, 1 optimum sink, 0 suboptimal sinks.



(b) $N = 15, k = 0.6$, 8 optimal sinks, 0 suboptimal sinks.



(c) $N = 15, k = 1.0$, 2 optimal sinks, 16 suboptimal sinks.




Fig. 2: Local optima networks (CM-LONs) for selected NPP instances with $N = 15$. Images are shown in 2D and a 3D projection (where the vertical dimension corresponds to fitness). Node sizes are proportional to their incoming strength, and edge thickness to their weight. Red nodes correspond to globally optimal sinks, while blue nodes to suboptimal sinks. For $k = 0.4$ the whole network is visualised, while for $k \in \{0.6, 1.0\}$, the fittest part of the network is shown.

collected from fitness landscapes and local optima networks [9,10]. Moreover, a new set of metrics can be gathered from computing the landscape sinks and funnel structure. We selected a subset of metrics after some preliminary experiments, including some that corroborate previous findings, and new local optima network metrics that intuitively relate to search dynamics. Figure 3 summarises the results, showing metrics for local optima network cardinality (1st row), neutrality (2nd row), sinks and funnels (3rd and 4th rows). The last row in Fig. 3 shows the empirical search cost of an Iterated Local Search (ILS) implementation, using a single bit-flip best-improvement hill-climber and a two bit-flip random perturbation.

Plot (a) confirms the surprising result, noted in previous studies [19,21], that the number of local optima remains virtually invariable across different values of k . Indeed most of the landscape metrics studied before: the size of the global and local basins, the correlation between basin size and fitness [21], and several barrier-tree metrics [19], are oblivious to the hardness phase transition. An exception is the number of global optima (plot b), which decreases progressively to 2 at about $k \approx 1$.

The rest of the landscape metrics reported in Fig. 3 can only be gathered using the local optima network model, specifically the compressed monotonic model (CM-LON) proposed in this article. The number of nodes in the CM-LON, gradually increases with increasing k , which correlates to (d) the number of distinct fitness values in the LON (Pearson’s correlation $r \approx 0.82$). Other metrics reflecting the amount of neutrality at the local optima network level are: (e) the mean number of nodes in a LON plateau and (f) the proportion of adjacent nodes that have the same fitness. Plot (e) reflects a sharp decrease from lower to higher k values, more noticeable for the largest N ; which suggests that the amount of neutrality is relevant to the phase transition.

Plot (g) presents the total number of sinks. This is divided into (h) the number of globally optimal sinks and (i) the number of suboptimal sinks. The bell shape for the number of globally optimal sinks appears because, for low values of k , global optima are part of a single LON plateau which gets compressed into one sink. Higher k show reduced neutrality, as seen in plots (d, e and f). From the NPP definition, higher values of k also mean that the number of global optima progressively decreases to reach 2 on average at $k \geq 1$. The theoretical minimum number of global optima is 2, where one solution is the negation of the other. An illustration of this phenomenon is provided in Figures 1 and 2. These metrics also hint to a transition starting to occur for values of k in the range 0.6 and 0.8. Again, a sharper change is observed for the largest N . The number of suboptimal sinks (plot i) is clearly relevant to search, as sinks act as traps for the search process. Once a suboptimal sink is reached it is not possible to escape, and the search stagnates in a suboptimal solution.

The plots in the fourth row show three metrics that also relate to search dynamics. Plot (j) reflects the average relative size of global optima funnels, that is, the fraction of local optima that lie on monotonic sequences leading to a global optimum sink. Clearly, the larger this value, the more chances a search

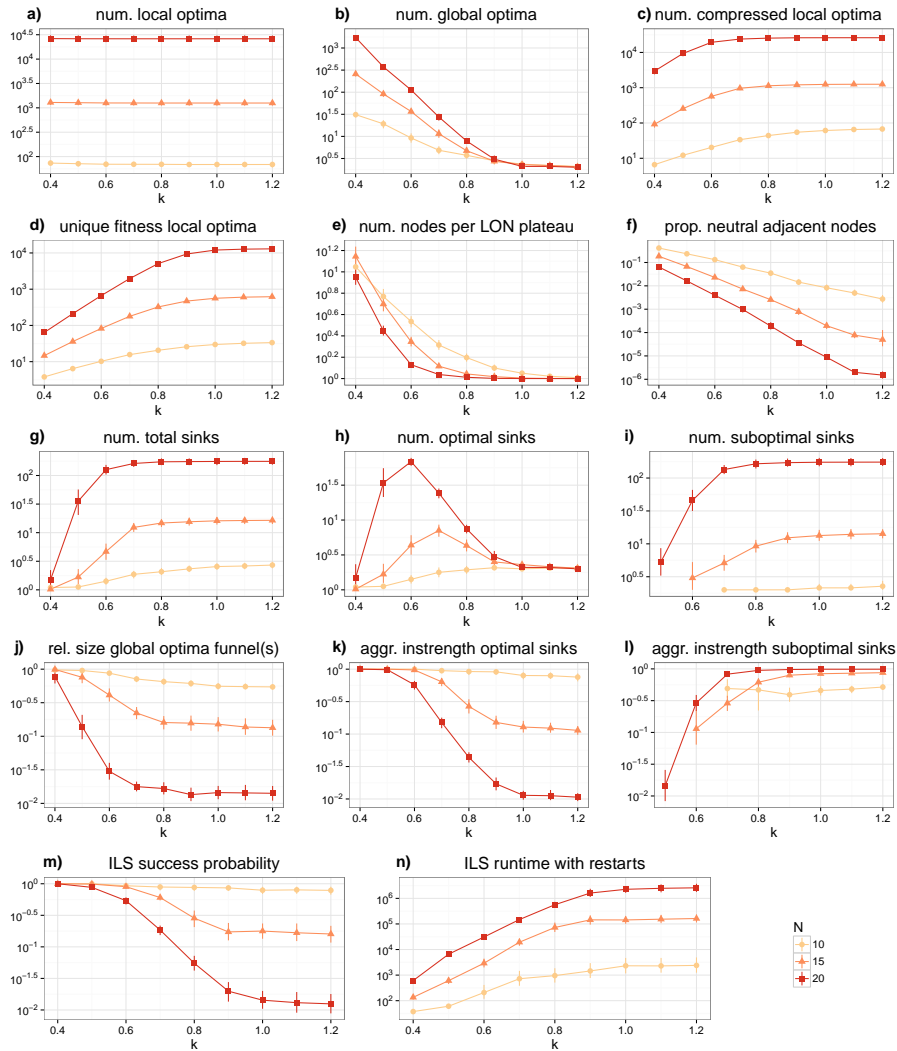


Fig 3: Local optima network features averaged over 30 instances per value of k (x-axis) and N (legend). The first row of plots illustrate cardinality metrics, the second describes neutrality, and the third statistics on the number of sinks. The fourth row shows metrics that relate to search dynamics: j) the fraction of local optima that lie on monotonic sequences to a global optimum sink, k) the aggregated incoming strength of optimal sinks, and l) the equivalent measure for suboptimal sinks. The last row shows the empirical search cost of an ILS algorithm on the same problem instances in terms of: m) success probability, and n) number of function evaluations with restarts. All plots are in semi-log scale.

process will have to find a path to an optimum. This metric decreases with k , more sharply for $N = 20$, with a transition between 0.6 and 0.8. Plots (k) and (l) report the weighted incoming degree (incoming strength) of the globally optimal and suboptimal sinks, respectively. These values are clearly relevant to search, the larger the aggregated incoming strength of optimal sinks, the higher the probability of a search process successfully reaching one of them. On the other hand, the larger the incoming strength of suboptimal sinks, the higher the chances of getting trapped. Again, a transition gradually occurs for values of k between 0.6 and 0.1.

Finally, the last row summarises the empirical cost of 1000 ILS runs in terms of (m) probability of success with a stopping condition consisting of 2^{15} function evaluations, and (n) number of function evaluations when the ILS is combined with random restarts until a global optimum is found [33]. We can notice a clear relationship between the aggregate incoming strength of globally optimal sinks and the probability of success. Figure 4a highlights this relationship with a scatter plot fitted with a univariate linear regression model. This confirms the strong correlation, with coefficient of determination $R^2 \approx 0.94$ (which corresponds to Pearson’s r correlation). Figure 4b is the scatter plot for the probability of success against the aggregate incoming strength of suboptimal sinks. Here, as we might expect, the relation is reversed. The number of suboptimal sinks is also inversely correlated to the probability of success ($r \approx -0.68$). These are useful relationships for future work on performance prediction since an approximation of the number of sinks and their incoming strength might be estimated using some sampling method.

Several of the metrics studied seem to reflect and explain the known hardness NPP phase transition. The transition, however, seems to appear earlier than the theoretical expected value of $k = 1.0$ [5,17,18], and is not exactly bracketed. However, these trends can be explained in the following way. In theory, the NPP computational phase transition becomes sharp only in the limit of infinite system size N and for $\log_2 M$, the number of bits in the input numbers, tending to infinity as well keeping k finite. In practice, we can only simulate finite systems because of computational limitations. A semi-rigorous argument [34] shows that in this case the transition point, k_c , becomes $k_c \approx 1 - \frac{\log_2 N}{2N}$ where the second term accounts for finite-size effects. For example, with $N = 15$, k_c is around 0.87 instead of one. This means that the phase transition is observed earlier and that it is not sharp, rather the system changes more gradually approaching it, and this is what we qualitatively observe in our numerical study.

5 Conclusions

Previous studies have failed to reveal clear links between the structure of fitness landscapes and the hardness phase transition known to exist on number partitioning problems when varying the critical parameter k . Most landscape metrics, except the number of global optima, are oblivious to the phase transition, and surprisingly remain invariable for easy and hard instances of this

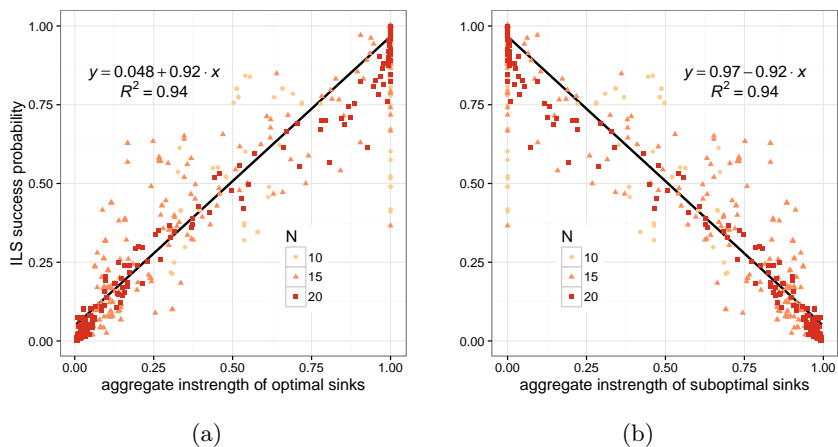


Fig. 4: ILS success probability against a) the aggregate incoming strength of the optimal sinks and b) the aggregate incoming strength of the suboptimal sinks. The lines are univariate linear regression models. The equation and coefficient of determination of each line are given in the plots.

problem. Our study sheds light into this puzzle, by considering new landscape metrics obtained from fully enumerated local optima networks. In particular we propose a local optima network model consistent with the *monotonic sequences* studied in theoretical chemistry, where the so-called multi-funnel structure of energy landscapes is well established. Our study reveals clear connections between the global structure of landscapes and the hardness phase transition. Easy instances show a dominant funnel structure leading to a set or connected global optima, or a small number or disjoint global optima (red nodes in Figures 2a and 2b). On the other hand, hard instances reveal multiple suboptimal funnels (blue nodes in Figure 2c), which explain why search gets trapped and is unable to escape with the commonly used perturbation operators. We found a strong correlation between the number, as well as the combined attracting strength, of suboptimal (blue) sinks and empirical search difficulty on the studied instances. Another important contribution of this work is to bring a more accessible visual approach to understanding search difficulty in combinatorial optimisation.

Future work will consider larger NPP instances using sampling, probe other number distributions, and most importantly, study whether other constraint satisfaction problems such as MAX-SAT reveal a similar global funnel structure explaining the hardness phase transition.

Acknowledgements. This work was supported by the Leverhulme Trust [award number RPG-2015-395] and by the UK's Engineering and Physical Sciences Research Council [grant number EP/J017515/1].

Data Access. All data generated during this research are openly available from the Stirling Online Repository for Research Data (<http://hdl.handle.net/11667/85>).

References

1. Garey, M.R., Johnson, D.S.: *Computers and Intractability*. Freeman, San Francisco, CA (1979)
2. Papadimitriou, C.H., Steiglitz, K.: *Combinatorial Optimization: Algorithms and Complexity*. Prentice-Hall, Englewood Cliffs, NJ (1982)
3. Gent, I.P., Walsh, T.: The SAT phase transition. In: *Proceedings of ECAI-96*. Volume 94., PITMAN (1994) 105–109
4. Culberson, J., Gent, I.P.: Frozen development in graph coloring. *Theoretical computer science* **265**(1) (2001) 227–264
5. Gent, I.P., Walsh, T.: Phase transitions and annealed theories: Number partitioning as a case study. In: *Proceedings of ECAI-96*, PITMAN (1996) 170–174
6. Gomes, C., Walsh, T.: Randomness and Structure. In Rossi, F., van Beek, P., Walsh, T., eds.: *Handbook of Constraint Programming*. Volume 2. Elsevier (2006) 639–664
7. Kambhampati, S.C., Liu, T.: Phase transition and network structure in realistic SAT problems. In: *Proceedings of the Twenty-Seventh AAAI Conference on Artificial Intelligence*. AAAI'13, AAAI Press (2013) 1619–1620
8. Martin, O.C., Monasson, R., Zecchina, R.: Statistical mechanics methods and phase transitions in optimization problems. *Theoretical computer science* **265**(1) (2001) 3–67
9. Tomassini, M., Vérel, S., Ochoa, G.: Complex-network analysis of combinatorial spaces: The NK landscape case. *Phys. Rev. E* **78**(6) (2008) 066114
10. Verel, S., Ochoa, G., Tomassini, M.: Local optima networks of NK landscapes with neutrality. *IEEE Transactions on Evolutionary Computation* **15**(6) (2011) 783–797
11. Doye, J.P.K., Miller, M.A., Wales, D.J.: The double-funnel energy landscape of the 38-atom Lennard-Jones cluster. *Journal of Chemical Physics* **110**(14) (1999) 6896–6906
12. Lunacek, M., Whitley, D., Sutton, A.M.: The impact of global structure on search. In Rudolph, G., Jansen, T., Beume, N., Lucas, S., Poloni, C., eds.: *Parallel Problem Solving from Nature - PPSN X*. Volume 5199 of LNCS., Berlin, Heidelberg, Springer Berlin Heidelberg (2008) 498–507
13. Kerschke, P., Preuss, M., Wessing, S., Trautmann, H.: Detecting funnel structures by means of exploratory landscape analysis. In: *Proceedings of the 2015 Annual Conference on Genetic and Evolutionary Computation*. GECCO '15, New York, NY, USA, ACM (2015) 265–272
14. Ochoa, G., Veerapen, N.: Deconstructing the Big Valley Search Space Hypothesis. In Chicano, F., Hu, B., García-Sánchez, P., eds.: *Evolutionary Computation in Combinatorial Optimization, EvoCOP 2016*. Volume 9595 of LNCS., Cham, Springer International Publishing (2016) 58–73
15. Ochoa, G., Veerapen, N.: Additional dimensions to the study of funnels in combinatorial landscapes. In: *Proceedings of the Genetic and Evolutionary Computation Conference 2016*. GECCO '16, New York, NY, USA, ACM (2016) 373–380
16. Herrmann, S., Ochoa, G., Rothlauf, F.: Communities of local optima as funnels in fitness landscapes. In: *Proceedings of the Genetic and Evolutionary Computation Conference 2016*. GECCO '16, New York, NY, USA, ACM (2016) 325–331

17. Ferreira, F.F., Fontanari, J.F.: Probabilistic analysis of the number partitioning problem. *Journal of Physics A: Mathematical and General* **31**(15) (1998) 3417
18. Mertens, S.: Phase transition in the number partitioning problem. *Phys. Rev. Lett.* **81**(20) (1998) 4281–4284
19. Stadler, P.F., Hordijk, W., Fontanari, J.F.: Phase transition and landscape statistics of the number partitioning problem. *Physical Review E* **67**(5) (2003) 056701
20. Flamm, C., Hofacker, I.L., P. F, S., Wolfinger, M.T.: Barrier trees of degenerate landscapes. *Zeitschrift für Physikalische Chemie International journal of research in physical chemistry and chemical physics* **216**(2/2002) (2002) 155–173
21. Alyahya, K., Rowe, J.E.: Phase transition and landscape properties of the number partitioning problem. In Blum, C., Ochoa, G., eds.: *Evolutionary Computation in Combinatorial Optimisation - 14th European Conference, EvoCOP 2014*. Volume 8600 of LNCS., Berlin, Heidelberg, Springer Berlin Heidelberg (2014) 206–217
22. Boese, K.D., Kahng, A.B., Muddu, S.: A new adaptive multi-start technique for combinatorial global optimizations. *Operations Research Letters* **16**(2) (1994) 101–113
23. Hains, D.R., Whitley, L.D., Howe, A.E.: Revisiting the big valley search space structure in the TSP. *Journal of the Operational Research Society* **62**(2) (2011) 305–312
24. Berry, R.S., Kunz, R.E.: Topography and dynamics of multidimensional interatomic potential surfaces. *Phys. Rev. Lett.* **74** (1995) 3951–3954
25. Wales, D.J.: Energy landscapes and properties of biomolecules. *Physical Biology* **2**(4) (2005) S86–S93
26. Becker, O.M., Karplus, M.: The topology of multidimensional potential energy surfaces: Theory and application to peptide structure and kinetics. *The Journal of Chemical Physics* **106**(4) (1997) 1495
27. Mézard, M., Mora, T., Zecchina, R.: Clustering of solutions in the random satisfiability problem. *Phys. Rev. Lett.* **94** (2005) 197205
28. Stadler, P.F.: Fitness landscapes. *Appl. Math. and Comput* **117** (2002) 187–207
29. Huynen, M.A., Stadler, P.F., Fontana, W.: Smoothness within ruggedness: the role of neutrality in adaptation. *Proc. Nat. Acad. Sci. U S A* **93**(1) (1996) 397–401
30. Barnett, L.: Ruggedness and neutrality - the NKp family of fitness landscapes. In Adami, C., Belew, R.K., Kitano, H., Taylor, C., eds.: *ALIFE VI, Proceedings of the Sixth International Conference on Artificial Life*, Cambridge, MA, USA, The MIT Press (1998) 18–27
31. Daolio, F., Verel, S., Ochoa, G., Tomassini, M.: Local Optima Networks of the Quadratic Assignment Problem. In: *2010 IEEE Congress on Evolutionary Computation (CEC)*. (2010) 1–8
32. Daolio, F., Tomassini, M., Vérel, S., Ochoa, G.: Communities of minima in local optima networks of combinatorial spaces. *Physica A: Statistical Mechanics and its Applications* **390**(9) (2011) 1684–1694
33. Auger, A., Hansen, N.: Performance evaluation of an advanced local search evolutionary algorithm. In: *Evolutionary Computation, 2005. The 2005 IEEE Congress on*. Volume 2., IEEE (2005) 1777–1784
34. Mertens, S.: The easiest hard problem: Number partitioning. In Percus, A., Istrate, G., Moore, C., eds.: *Computational Complexity and Statistical Physics*. Volume 125 of *The Santa Fe Institute Studies in the Sciences of Complexity*. Oxford University Press, New York, NY, USA (2006) 125–139