

Perturbation Strength and the Global Structure of QAP Fitness Landscapes

Gabriela Ochoa¹ and Sebastian Herrmann²

¹ University of Stirling, Scotland, United Kingdom,
gabriela.ochoa@stir.ac.uk

² Hochschule RheinMain, Wiesbaden, Germany,
sebastian.herrmann@hs-rm.de

Abstract. We study the effect of increasing the perturbation strength on the global structure of QAP fitness landscape induced by Iterated Local Search (ILS). The global structure is captured with Local Optima Networks. Our analysis concentrates on the number, characteristics and distribution of funnels in the landscape, and how they change with increasing perturbation strengths. Well-known QAP instance types are considered. Our results confirm the multi-funnel structure of QAP fitness landscapes and clearly explain, visually and quantitatively, why ILS with large perturbation strengths produces better results. Moreover, we found striking differences between randomly generated and real-world instances, which warns about using synthetic benchmarks for (manual or automatic) algorithm design and tuning.

Keywords: Local Optima Network, Quadratic Assignment Problem, QAP, Iterated Local Search, Perturbation Strength, Fitness Landscapes

1 Introduction

The Quadratic Assignment Problem (QAP) requires the assignment at minimal cost of a set of facilities to a set of locations, given the flows between facilities and the distances between locations. The QAP is one of the most difficult combinatorial optimisation problems. Current exact algorithms can solve mostly problems of up to 30 to 40 facilities, therefore, metaheuristics are frequently used to solve larger instances. The most successful are Hybrid Evolutionary Algorithms [1, 2] and Iterated Local Search (ILS) with variable perturbation strengths [3, 4].

Several studies have analysed QAP fitness landscapes [1, 3, 5]. The local properties are usually studied through an autocorrelation analysis, while the global structure through a fitness distance correlation analysis. The existence of plateaus and the structure of basins have also been studied [5]. These studies suggest that QAP instances have *unstructured* landscapes. The distances between local optima and the best-known solutions, as well as the average distances between local optima are very close to the landscape diameter (n the instance size). The local optima are neither restricted to a small region of the search space, nor seem to be correlated, QAP landscapes do not seem to have a ‘big-valley’.

Beyond fitness distance correlation analysis, there are no tools available to understand the global structure of fitness landscapes. Local Optima Networks (LONs) [6], fill this gap by providing a compressed model of landscapes, where nodes are local optima, and edges possible transitions among them. They model the distribution and connectivity pattern of local optima, and thus help to characterise the underlying landscape global structure. LONs have been recently used to study the multi-funnel structure of fitness landscapes [7–9]. A funnel refers to a grouping of local optima, forming a coarse-level gradient towards a low cost solution at the end. When sub-optimal funnels exist, search can get trapped and thus fail to reach the global optimum despite a large computing time.

Iterated local search is a simple yet powerful search strategy. It works by iteratively alternating an intensification stage (local search) with a diversification stage (perturbation or kick). We will refer to the strength of a perturbation as the number of solution components that are modified. Such components are, for example, the number of jobs to move in production scheduling, or the number of edges to interchange in the Travelling Salesperson Problem (TSP) [10]. Recent studies using LONs have shown that increasing the perturbation strength of ILS can ‘smooth’ the multi-funnel structure, i.e., cause some funnels to disappear or merge, with the effect of improving the algorithm performance [11, 12].

For some problems, such as the TSP [10, 12], an appropriate perturbation strength can be small and rather independent of the instance size. ILS implementations to solve the QAP, instead, have shown to benefit from large perturbation sizes. As reported in [10], the best perturbation size depends on the particular instance. For some instances, altering as many as 75% of the solution components produced the best performance. We argue that this behaviour can be better understood by studying the underlying landscape global structure.

The effect of increasing the perturbation strength on the global structure of QAP fitness landscapes has not yet been studied. In this article, we use LONs to characterise and contrast the landscapes of QAP instances of different classes and perturbation strengths. Our goal is twofold. First, to offer a deeper understanding of why an increased perturbation strength proves advantageous for the QAP. Second, to illustrate visually and quantitatively, the structural differences between real-world and synthetic QAP instances. Another contribution of this article is a more rigorous description of the notion of funnels, formalising the notion of *monotonic sequences* from the study of energy landscapes in theoretical chemistry [13] to the context of LONs for combinatorial optimisation.

2 The Quadratic Assignment Problem

A solution to the QAP is generally written as a permutation s of the set $\{1, 2, \dots, n\}$, where s_i gives the location of item i . Therefore, the search space is of size $n!$. The cost, or fitness function associated with a permutation s is a quadratic function of the distances between the locations, and the flow between the facilities, $f(s) = \sum_{i=1}^n \sum_{j=1}^n a_{ij} b_{s_i s_j}$, where n denotes the number of facili-

ties/locations and $A = \{a_{ij}\}$ and $B = \{b_{ij}\}$ are the distance and flow matrices, respectively.

Our goal is to visualise and characterise in detail the global structure of QAP fitness landscapes. Therefore, we consider a group of 8 instances from the QAPLIB³ [14] with sizes ranging from 30 to 42 facilities, which are of moderate size yet not trivial to solve. Specifically, we selected the largest available real-world instances around this range, and complemented them with instances of similar sizes from the other types. According to [15, 3], most QAPLIB instances can be classified into four types. We selected two instances of each type as indicated below.

1. *Uniform random distances and flows.* In these problems, denoted by *tainma*, where nn is the problem size, flows and distances are randomly drawn from a uniform distribution. They are known to be the hardest to solve optimally, however, heuristic methods generally find solutions 1 or 2 per cent above the optimum in short computation time [15]. We selected two instances of this group: **tai30a** and **tai35a**.
2. *Random flows on grids.* These problems consider a rectangular tiling of $n_1 \times n_2$ squares of unit size. A location is one of these squares and the distances between them are measured. The flows are randomly generated, but not necessarily uniformly. These problems are known to be symmetrical and to have multiple of 4 ($n_1 \neq n_2$) or 8 ($n_1 = n_2$) different optimal solutions. We selected two instances of this group: **nug30** and **sko42**.
3. *Real-world problems.* These problems arise from practical applications. We briefly describe the instances used in this article. The Krarup and Pruzan instances, denoted by *kra*, contain real world data and were used to plan the Klinikum Regensburg in Germany. In the Steinberg's instances, denoted by *ste*, the goal is to minimise the length of connections between units that have to be placed on a rectangular grid [14]. We selected two instances of this group: **kra30a** and **ste36a** (notice that there are no real-world instances of size around 40 in QAPLIB).
4. *Random real-world like problems.* These instances, denoted by *tainnb*, are randomly generated in a way that they resemble the structure of the real-world instances. We selected two instances of this group: **tai30b** and **tai40b**.

3 Algorithms and Definitions

We use the implementation of ILS by Stützle [3], which allows us to explore the effect of different perturbation strengths. The local search stage uses a first improvement hill-climbing variant with the pairwise (2-exchange) neighbourhood. This operator swaps any two positions in a permutation. The perturbation operator exchanges k randomly chosen items, which corresponds to a random move in the k -opt neighbourhood. Algorithm 1 outlines the pseudo-code.

³ <http://www.seas.upenn.edu/qaplib/>

Algorithm 1 Iterated Local Search (ILS)

Require: Search space S , Fitness function $f(S)$,
Perturbation strength k , Stopping threshold t

- 1: Choose initial random solution $s_0 \in S$
- 2: $l \leftarrow \text{LocalSearch}(s_0)$
- 3: $i \leftarrow 0$
- 4: **repeat**
- 5: $s' \leftarrow \text{Perturbation}(l, k)$
- 6: $l' \leftarrow \text{LocalSearch}(s')$
- 7: **if** $f(l') \leq f(l)$ **then**
- 8: $l \leftarrow l'$
- 9: $i \leftarrow 0$
- 10: **end if**
- 11: $i \leftarrow i + 1$
- 12: **until** $i \geq t$
- 13: **return** l

3.1 Monotonic LON Model

Local optima. We assume a search space S with a fitness function $f(S)$ and a neighbourhood function $N(s)$. A local optimum, which in the QAP is a minimum, is a solution l such that $\forall s \in N(l), f(l) \leq f(s)$. Notice that the inequality is not strict, in order to allow the treatment of neutrality (local optima of equal fitness), which we found to occur in some QAP instances. The set of local optima, which corresponds to the set of nodes in the network model, is denoted by L .

Monotonic perturbation edges. Edges are directed and based on the perturbation operator (k -exchange, $k > 2$). There is an edge from local optimum l_1 to local optimum l_2 , if l_2 can be obtained after applying a random perturbation (k -exchange) to l_1 followed by local search, and $f(l_2) \leq f(l_1)$. These edges are called *monotonic* as they record only non-deteriorating transitions between optima. Edges are weighted with estimated frequencies of transition, specifically, integers indicating the number of times an edge was visited during the sampling process (see Section 4). The set of edges is denoted by E .

Monotonic LON. Is the weighted, directed graph $MLON = (L, E)$ where nodes are the local optima L , and edges E are the monotonic perturbation edges.

3.2 Compressed Monotonic LON Model

We have observed neutrality at the LON level (i.e. connected sets of optima that share the same fitness value) on several combinatorial problems. This lead us to propose a coarser LON model [9], which compresses the local optima that are connected by neutrality into single nodes. We found this type of neutrality on some QAP instances, specifically, the grid-based and real-world instances (see Section 2). Therefore, we considered the compressed model in this study.

Compressed local optima. A compressed local optimum is a set of connected nodes in the MLON with the same fitness value. Two nodes in the MLON are connected if there is a monotonic perturbation edge between them. The set of connected MLON optima with the same fitness, denoted by CL , corresponds to the set of nodes in the Compressed Monotonic LON model.

Compressed Monotonic LON. Is the weighted, directed graph $CMLON = (CL, CE)$, where nodes are the compressed local optima CL . The edges CE are aggregated from the monotonic edge set E by summing up the edge weights.

Monotonic Sequence. A monotonic sequence is a path of connected nodes $MS = \{cl_1, cl_2, \dots, cl_s\}$ where $cl_i \in CL$. By definition of the edges, $f(cl_i) \leq f(cl_{i-1})$. There is a natural end to every monotonic sequence, cl_s , when no improving transitions can be found. In the directed CMLON network, cl_s will be a node without outgoing edges (or sink)⁴.

Funnel. A funnel can be loosely described a grouping of local optima, conforming a coarse-grained gradient towards a low cost optimum. More formally, we characterise funnels in the CMLON as the aggregation of all monotonic sequences ending at the same point (or sink). Funnels can be seen as basins of attraction at the level of local optima.

4 Experimental Setting

Exhaustive enumeration of the search space, and thus complete extraction of the LON models, is not feasible for permutation sizes larger than 10. Our instances have sizes $n \geq 30$, therefore sampling is required. The sampling procedure consists of running an instrumented version of Stützle [3] ILS (described in Section 3), where the stopping threshold is set to $t = 10\,000$ iterations without any improvement. This serves the purpose of empirically estimating the funnel ends, i.e., solutions at the end of the ILS trajectory, where escaping is difficult, if not impossible. Our ILS only accepts improving or equal fitness from perturbation moves in order to model monotonic sequences. While running ILS, we store in a set L all the unique optima obtained after the local search stage, and in a set E all the unique edges obtained after a perturbation followed by local search. To construct the MLONs for each instance and perturbation strength, these sets of nodes and edges are aggregated over 200 runs, started from different random configurations. Five perturbation strengths p , corresponding to k -exchanges, are explored (summarised in Table 1), ranging from $k = \frac{n}{8}$ to a complete restart. Once the MLONs are constructed, we proceed to identify the connected components with shared fitness, and thus construct the respective compressed models.

There are multiple performance and network metrics that can be computed and used to understand search difficulty and landscape structure. We selected a

⁴ In directed graphs, a node without outgoing edges is called a *sink*.

Table 1: Perturbation strengths.

flag (p)	1	2	3	4	5
size (k)	$\frac{n}{8}$	$\frac{n}{4}$	$\frac{n}{2}$	$\frac{3n}{4}$	<i>restart</i>

minimal subset characterising the key global structural properties. ILS performance is measured with two metrics: (i) success rate, i.e., the proportion of runs that reached a global optimum, and (ii) the normalised proximity to the global optimum evaluation (i.e the inverse of the performance gap). These metrics are gathered from 200 independent runs on each instance and perturbation strength.

Table 2: Performance and local optima network metrics.

Perf.	<i>success</i>	Success rate of finding a global optimum
Metrics	<i>proximity</i>	Normalised proximity to the global optimum evaluation
	<i>noptima</i>	Number of optima (including both local and global)
	<i>compression</i>	Ratio of number of compressed to total number of optima
Network Metrics	<i>pglobal</i>	Proportion of funnels that are globally optimal
	<i>strength</i>	Normalised incoming strength of globally optimal sink(s)

The landscapes’ global structure is assessed using four characteristics, summarised in Table 2. The number of optima (*noptima*) is the number of nodes of the MLON model. Since we compute the compressed model CMLON, we report also the ratio of the number of compressed optima to the number of optima (*compression*). For characterising the multi-funnel structure, we measure the proportion of funnels that end at the global optimum level (*pglobal*). This is the ratio of the number of global funnels to the total number of funnels. Finally, as a measure of the centrality and reachability of the global optima, we compute the normalised incoming strength (weighted degree) of the global optimal sinks (*strength*). This is computed as the sum of the incoming strengths of the globally optimal sinks divided by the sum of the incoming strengths for all sinks.

5 Results

5.1 Visualisation

A first step in analysing the structure of networks is often to visualise them. Figure 1 illustrates four compressed MLONs of a single QAP instance, **sko42**

at different perturbation strengths $p \in \{1, 3, 4, 5\}$. **sko42** is representative for other QAP instances and also the largest in our set. The number of optimal and suboptimal funnels and the ILS success rate are also indicated. Each node is a compressed optimum, and edges are monotonic transitions with the corresponding perturbation strength.





Fig. 1: Compressed monotonic local optima networks for instance sko42 at four different perturbation strengths $p \in \{1, 3, 4, 5\}$. The number of funnels f (*optimal, suboptimal*) and the ILS success rate sr are indicated. The size of nodes is proportional to their incoming strength. Red highlights the global optimal funnels, while blue suboptimal funnels. Funnel ends (sinks) are indicated in more intense colours.

The networks in Figure 1 capture the whole set of sampled nodes and edges for each network, whose sizes range from 1,976 nodes (plot (d)) to 3,234 nodes (plot (a)). Plots were produced with the R statistical language using *force-*


directed layout methods as implemented in the `igraph` library [16]. The decorations reflect features relevant to search dynamic; the size of nodes is proportional to their incoming weighted degree (strength), which indicates how much a node ‘attracts’ the search process. Red nodes belong to the funnel(s) containing a global optimum, while blue nodes to suboptimal funnels. The funnels’ terminating nodes (sinks) are highlighted with a black outline and a more intense colour. A visual inspection reveals clear structural differences among the four perturbation strengths. For the smallest strength ($p = 1$, plot (a)), four global optima can be seen in red, but there is no clear funnel basin structure. However, for $p = 2$ (plot (b)), a clear pattern emerges, showing the basins (i.e collection of monotonic sequences with a common ending point) of four optimal funnels in red, and a number of suboptimal funnels in blue. The success rate increases to 0.53, the highest for this instance. For a stronger perturbation, ($p = 4$, plot (c)) the performance deteriorates. The four global optimal funnels can still be seen, but their basins become small, and a multitude of sub-optimal funnels appear. Finally, for a complete restart ($p = 5$, plot (d)), only one of the four global optima was found by the sampling process, the number of sub-optimal funnels is almost 200 and the success rate is close to zero.

In order to have a view of the global structure of the remaining three types of QAP instances, Figure 2 shows the network plots for a uniform random instance tai30a (plots (a) and (b)), a real-world instance kra30a (plots (c) and (d)), and a random real-world like instance tai40a (plots (e) and (f)). Due to space constraints, only two perturbation strengths are shown, $p = 1$ (left plots) and the value of p producing the highest success rate (right plots). The most striking observation from these plots are the structural differences between the randomly generated instances and the real-world instance. The random instances have a single global optimum, whereas the real-world (and indeed the grid based instances, see Fig. 1), generally have several global optima. The random uniform instances (tai30a, plots (a) and (b)) are very difficult to solve to optimality. Even for the best-performing perturbation strength ($p = 3$, plot (b)), the single global optimum belongs to a tiny funnel, so it is difficult to find. On the other extreme, the random real-world like instances (e.g. tai40b, plots (e) and (f)), are far too easy to solve, reaching 100% success for $p = 3$. Importantly, for both perturbation strengths, a single funnel structure is observed (there are no blue nodes in plots (e) and (f)), which clearly contrasts with the multi-funnel nature of the other instances.


The tai40b networks (plots (e) and (f)) show some intermediate nodes with high strengths (larger pink nodes) in the path to the global optima (red node with black outline). However, there are escape edges from these intermediate attractors, so the perturbation operator is able to overcome them and reach the global optimum. These instances are therefore easy to solve, and the success rate will be 1.0 even for a low perturbation, given enough running time. Regarding the real-world instance kra30a, for a perturbation strength $p = 3$ (plot (d)), the four optimal funnels observed for $p = 1$ (plot (c)) merge into a single funnel with a large incoming strength, explaining the high success rate in this case.




(a) tai30a, $p=1$, $f=(1,188)$, $sr=0.015$




(b) tai30a, $p=3$, $f=(1,94)$, $sr=0.03$




(c) kra30a, $p=1$, $f=(4,63)$, $sr=0.18$



(d) kra30a, $p=4$, $f=(1,0)$, $sr=0.97$



(e) tai40b, $p=1$, $f=(1,0)$, $sr=0.48$



(f) tai40b, $p=3$, $f=(1,0)$, $sr=1.0$

Fig. 2: Compressed monotonic local optima networks for instances of the three remaining types and two perturbation strengths, $p = 1$ (left) and the best performing perturbation (right). For more details see caption of Fig. 1

5.2 Structural and performance metrics

Figure 3 shows the network and performance metrics described in Table 2 for the eight QAP instances selected. Results were collected for five perturbation strengths as indicated in Table 1. For most instances, the best performance, as measured by both success rate and proximity to the optimum cost, is achieved with an intermediate perturbation strength of $p = 3$, which corresponds to 50% changes to the solution. A perturbation strength of $p = 4$ (75% alterations to the solution), produces the best performance for one of the instances, kra30a. For the random uniform instances (tai30 and tai35), success rates are rather low, however, the solutions found are not far in evaluation from the global optimum as indicated by the proximity metric. The network metrics, *strength* and *pglobal*, seem to correlate well with the performance metrics. The reduction of the number of sub-optimal funnels and the increased strength of the global optimal sink, obtained with perturbation strengths $p = \{3, 4\}$, offer an explanation for the increased success rate from the point of view of the landscape global structure.

The most remarkable observations from the plots in Figure 3 are the notable differences between the randomly generated instances and the real-world and grid-based instances. Looking at the *compression* metric, we can observe that both the random uniform and random real-world like instances show zero compression, which means that there is no neutrality at the level of local optima. The grid-based (nug30 and sko42) and real-world (kra30a and sko42) instances, on the other hand, show neutrality at the local optima level, which decreases with the perturbation strength, indicating that local optima with the same fitness become connected and thus merged. The uniform random instances (tai30a and tai35a) have very low levels of *strength* and ratio of global optimal funnels (*pglobal*), whereas the random real-world like instances (tai30b and tai40b) go to the other extreme, showing a single funnel leading to the global optimum.

6 Conclusion

We have extracted and analysed the compressed monotonic local optima networks induced by iterated local search with different perturbation strengths on QAP instances of different types. Our results confirm that large perturbation strengths, of around 50% or occasionally even 75% alterations to the solution, produce the best performance. Our analysis explains this behaviour; a multi-funnel structure generally occurs with low perturbation strengths. With increased perturbation strength, the number of sub-optimal funnel decreases, while the size of the optimal funnel increases, which facilitates reaching the global optimum. We found striking differences between the randomly generated instances, the real-world instances and the grid-based instances. The latter have more than one global optimum, and several local optima sharing the same fitness, which is not the case on the random instances. Moreover, the uniform random instances are very difficult to solve to optimality; they show multiple funnels with very small basins, which do not merge with increased perturbation. On the other extreme, the random real-world like instances are very easy to solve

to optimality; their global structure shows a single funnel with a large basin. This is in stark contrast with the multi-funnel structure of the real-world and grid-based instances.




Fig. 3: Structural and performance metrics (as defined in Table 2) for five perturbation strengths (as defined in Table 1).

We argue that care should be taken when using randomly generated instances to improve the manual or automatic design of heuristic methods. It is not clear that knowledge extracted from random instances will generalise to real-world instances. Future work will explore the effect of adding crossover, and whether the most effective perturbation strength can be inferred from cheap estimations of the global structure of the underlying landscapes.

Acknowledgements. Thanks are due to Thomas Stützle who shared with us his QAP Iterated Local Search source code.

References

1. Merz, P., Freisleben, B.: Fitness landscape analysis and memetic algorithms for the quadratic assignment problem. *IEEE Trans. Evolutionary Computation* 4(4), 337–352 (2000)
2. Drezner, Z., Misevicius, A.: Enhancing the performance of hybrid genetic algorithms by differential improvement. *Computers and Operations Research* 40(4), 1038 – 1046 (2013)
3. Stützle, T.: Iterated local search for the quadratic assignment problem. *European Journal of Operational Research* 174(3), 1519–1539 (2006)
4. Benlic, U., Hao, J.K.: Breakout local search for the quadratic assignment problem. *Applied Mathematics and Computation* 219(9), 4800–4815 (2013)
5. Tayarani-N., M.H., Prügel-Bennett, A.: Quadratic assignment problem: a landscape analysis. *Evolutionary Intelligence* 8(4), 165–184 (2015)
6. Ochoa, G., Tomassini, M., Verel, S., Darabos, C.: A study of nk landscapes’ basins and local optima networks. In: *Genetic and Evolutionary Computation Conference, GECCO*. pp. 555–562. ACM (2008)
7. Herrmann, S., Ochoa, G., Rothlauf, F.: Communities of local optima as funnels in fitness landscapes. In: *Genetic and Evolutionary Computation Conference GECCO*. pp. 325–331 (2016)
8. Ochoa, G., Veerapen, N.: Mapping the global structure of TSP fitness landscapes. *Journal of Heuristics* pp. 1–30 (2017)
9. Ochoa, G., Veerapen, N., Daolio, F., Tomassini, M.: Understanding Phase Transitions with Local Optima Networks: Number Partitioning as a Case Study. In: *Evolutionary Computation in Combinatorial Optimization, (EVOCOP)*. LNCS, vol. 10197, pp. 233–248. Springer (2017)
10. Lourenço, H.R., Martin, O.C., Stützle, T.: *Iterated Local Search: Framework and Applications*, pp. 363–397. Springer US, Boston, MA (2010)
11. Herrmann, S., Herrmann, M., Ochoa, G., Rothlauf, F.: Shaping communities of local optima by perturbation strength. In: *Genetic and Evolutionary Computation Conference, GECCO*. pp. 266–273 (2017)
12. McMenemy, P., Veerapen, N., Ochoa, G.: How erturbation Strength Shapes the Global Structure of TSP Fitness Landscapes. In: *Evolutionary Computation in Combinatorial Optimization, (EVOCOP)*. p. to appear. LNCS, Springer (2018)
13. Berry, R.S., Kunz, R.E.: Topography and dynamics of multidimensional interatomic potential surfaces. *Phys. Rev. Lett.* 74, 3951–3954 (1995)
14. Burkard, R.E., Karisch, S.E., Rendl, F.: QAPLIB – a quadratic assignment problem library. *Journal of Global Optimization* 10(4), 391–403 (1997)
15. Taillard, E.: Comparison of iterative searches for the quadratic assignment problem. *Location Science* 3(2), 87–105 (1995)
16. Csardi, G., Nepusz, T.: The igraph software package for complex network research. *InterJournal Complex Systems*, 1695 (2006)