ILS-perturbation based on local optima structure for the QAP problem

Everardo Gutiérrez and Carlos A. Brizuela

Computer Science Department, CICESE Research Center
Km 107 Carr. Tijuana-Ensenada, Ensenada, B.C., México
{egutierr, cbrizuel}@cicese.mx, +52–646–1750500

Abstract. Many problems in AI can be stated as search problems and most of them are very complex to solve. One alternative for these problems are local search methods that have been widely used for tackling difficult optimization problems for which we do not know algorithms which can solve every instance to optimality in a reasonable amount of time. One of the most popular methods is what is known as iterated local search (ILS), which samples the set of local optima searching for a better solution. This algorithm's behavior is achieved by some mechanisms like perturbation which is a key aspect to consider, since it allows the algorithm to reach a new solution from the set of local optima by escaping from the previous local optimum basis of attraction. In order to design a good perturbation method we need to analyze the local optima structure such that ILS leads to a good biased sampling. In this paper, the local optima structure of the Quadratic Assignment Problem, an NP-hard optimization problem, is used to determine the required perturbation size in the ILS algorithm. The analysis is focused on verifying if the set of local optima has the "Big Valley (BV)" structure, and on how close local optima are in relation to problem size. Experimental results show that a small perturbation seems appropriate for instances having the BV structure, and for instances having a low distance among local optima, even if they do not have a clear BV structure. Finally, as the local optima structure moves away from BV a larger perturbation is needed.
Keywords: Iterated Local Search, Big Valley, Perturbation Length, Quadratic Assignment Problem.

1 Introduction

Most problems in AI can be classified as general search problems and many methods have been proposed to deal with them [17]. However, the existence of problems for which there are no known algorithms which can ensure optimality in a reasonable amount of time [6] has motivated the development of alternative methods to obtain acceptable solutions from a practical point of view. One of the alternatives is the method known as *local search* (LS), which has been successfully used in practice for difficult combinatorial optimization problems [1, 5, 6].

Local search algorithms (LSA) have a general behavior based on the following idea: take an initial solution and modify it until no further improvements are possible. These algorithms need to define a *neighborhood* function N which represents a map $N : S \rightarrow 2^S$, such that it defines for each solution s in the set of all feasible solutions S a subset $N(s) \subset S$ of *neighbors* of s [1]. This function defines the structure over which the search must be done, and this structure is called *search graph* [4], *fitness landscape* [16, 19], or *state space* [17].

There are a lot of suggested algorithms based on this behavior and Iterated Local Search (ILS) is one of them [13, 11]. ILS explores the set of local optima applying a local search and a perturbation mechanism to restart the search. While a better local search can reach a better local optimum by itself, this may make it difficult to perturb this solution such that a different local optimum can be reached. Many perturbation methods have been proposed and they are one of the key aspects to consider in the design of ILS algorithms [11].

The Quadratic Assignment Problem (QAP) is a very important combinatorial optimization problem from a practical and theoretical point of view. Since QAP belongs to *NP-hard* class [6], many alternative methods such as LS, have been proposed to deal with it. In this paper we analyze the local optima structure of an instance set, and its relation with the perturbation mechanism in a simple ILS algorithm.

The remainder of this paper is organized as follows. Section 2 explains the idea of ILS and their mechanisms. In section 3 we state the QAP problem. Section 4 gives the concepts of "Global Convexity" and "Big Valley", and shows some QAP local optimum structures. Section 5 presents the experimental setup and their results. Finally, section 6 gives the conclusions and ideas for future research.

2 Iterated Local Search (ILS)

Iterated Local Search is among the best performing local search algorithms for some well known combinatorial optimization problems [9, 18, 11]. The essence of ILS is: one iteratively builds a sequence of solutions generated by an embedded heuristic [11]. Let f be the cost function of our combinatorial optimization problem; f is to be minimized. Let S be the set of all solutions s. Finally, the local search procedure $LocalSearch$ defines a mapping from the set S to the smaller set S^* of locally optimal solutions s^*, $LocalSearch$ procedure samples S^*. Then, given the mapping from S to S^*, we need to reduce the costs found without modifying $LocalSearch$. ILS tries to avoid the disadvantages of random restart by exploring S^* using a walk that steps from one s^* to a "nearby" one, and can be described with the following high level steps:

procedure *Iterated Local Search*
    s_0 = **GenerateInitialSolution()**
    s^* = **LocalSearch(**s_0**)**
    repeat
        s' = **Perturbation(**s^*, *history***)**
        $s^{*'}$ = **LocalSearch(**s'**)**

$s^* = $ **AcceptanceCriterion**(s^*, $s^{*'}$, *history*)
until termination condition met

end

According to this general architecture, ILS behavior is defined by the routines: **GenerateInitialSolution**, **LocalSearch**, **Perturbation** and **AcceptanceCriterion**. The main aspects of each one of them are briefly described in the following:

Initial solution. The starting point s_0 can be generated by different methods (random, greedy starting, etc.) and it only has influence on which local optimum is going to be visited first.

Local search. Local search is usually viewed as a black box which allows us to sample the set S^*. For some problems, as in case of TSP, the better the local search, the better the corresponding ILS [11]. However, an excellent local search should systematically undo the work done by the perturbation.

Perturbation. The perturbation should allow the algorithm to reach a solution which can be a seed for a new local optimum using the local search procedure. Then, the perturbation should be large enough to "escape" from the current local minimum basis of attraction, but not too large to give an almost random solution. Previous work claim that while for TSP a perturbation of fixed/small size (double-bridge) has shown a good performance, for QAP a relatively large perturbation size seems to be necessary [11].

Acceptance criterion. This procedure determines whether $s^{*'}$ is accepted or not as the updated current solution. Together with *Perturbation*, it controls the balance between intensification and diversification of the search over S^*.

We will define each ILS routine such that the resulting algorithm is going to be used to tackle QAP, a well known combinatorial optimization problem which is stated in the next section.

3 Quadratic Assignment Problem

The Quadratic Assignment Problem is a combinatorial optimization problem and can be described as the problem of assigning a set of facilities to a set of locations with given distances between locations and flows between facilities. The goal is to find the better assignment of facilities to the locations such that the sum of the product between flows and distances is minimal. Many practical problems can be formulated as QAP and it has been proven that the decision version of this problem is *NP-complete* [6].

Given n facilities and n locations, two n x n matrices A and B for distances and flows respectively, the QAP can be stated as follows:

$$\min_{\phi \in \Phi} \sum_{i=1}^{n} \sum_{j=1}^{n} b_{ij} a_{\phi_i \phi_j} \tag{1}$$

Table 1. Some QAPLIB instances. Problem size, best known values, and lower bounds.

Name	N	Best	Bound
bur26a	26	5426670	5426670
bur26b	26	3817852	3817852
bur26c	26	5426795	5426795
bur26d	26	3821225	3821225
chr25a	25	3796	3796
esc32a	32	130	103
esc32b	32	168	132
kra30a	30	88900	88900
lipa20a	20	3683	3683
lipa60a	60	107218	107218
lipa90a	90	360630	360630
tai30a	30	1818146	1529135
tai30b	30	637117113	589470167
tai60a	60	7205962	5555095
tai60b	60	608215054	50113782

where Φ is the set of all permutations of n numbers, ϕ_i gives the location of facility i and $b_{ij}a_{\phi_i\phi_j}$ is the cost contribution of assigning facility i to location ϕ_i and facility j to location ϕ_j.

Table 1 shows some of the QAPLIB [1] instances. **N** is the problem size (number of facilities and locations), **Best** is the best known solution and **Bound** is the best known lower bound.

We will analyze in next section the local optima structure of these instances in order to choose a good perturbation. The analysis is focused on verifying if the cost-distance relation in the set of local optima shows a BV structure.

4 Global Convexity

The correlation of solution's quality (cost, fitness, etc.) and distance between solutions has been studied in many ways [2, 10, 12, 14, 19]. The existence of correlation would mean that the search graph has some structure which could be used to guide the search in order to obtain better results. That is, the correlation could be used as a tool in the design of search strategies which take advantage of the search graph structure.

Boese *et al.* [3] analyze the relationships among local optima for two problems: TSP and Graph Bisection. They analyzed a search graph for each problem generating a set of local optima and measuring the distance among them. For the analyzed search graphs, the better solutions have a shorter mean distance to the others. These results suggest the existence of a globally convex structure [8] in the set of local optima, which they refer to as the "Big Valley" structure.

[1] http://www.seas.upenn.edu/qaplib/

Reeves [15] found the same structure for the Flow Shop Problem, using different distance measures and neighborhood structures. Figure 1 shows an example of a function that is "globally convex".

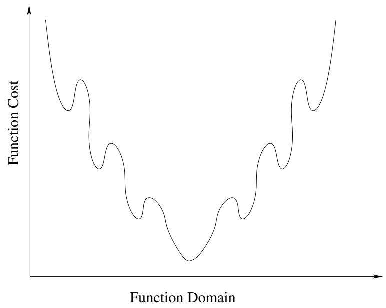

Fig. 1. Example of a globally convex function.

Given a neighborhood N for a problem Π and its respective search graph G, let S^* be the set of local optima with respect to N. In order to verify whether the search graph has the global convexity structure we must analyze the set S^*. Whenever it is not possible to obtain all $s^* \in S^*$ we can obtain a representative sample and measure solutions' quality and distance among them.

5 Experimental Results

5.1 Cost-Distance correlation

Here we compute the correlation of the cost and the mean distance between a set of local optima over a neighborhood that swap the location of two facilities. For each instance, a local search was performed 6400 times, and the final solution (a local optimum) was taken. The general procedure was a simple greedy local search which explores the neighborhood by searching for a better solution. If there is no better solution the algorithm stops and returns the final solution, which is, by definition, a local optimum.

Instances from QAPLIB show different cost-distance structures among their local optima even if they have the same input distributions. Figures 2, 3, 4, 5, and 6 show the cost-distance structures found for these instances.

We can see that some instances, as those shown in Figure 2, have structures very different from BV, while others like chr25a in Figure 3 and tai30a in Figure 6, show a shape very close to BV. However, it is not easy to distinguish some structures only by visual inspection of these graphs.

We will try to relate these structures with the experimental results produced by the ILS, and whenever this relation is not clear we will use other measures to explain the ILS behavior.

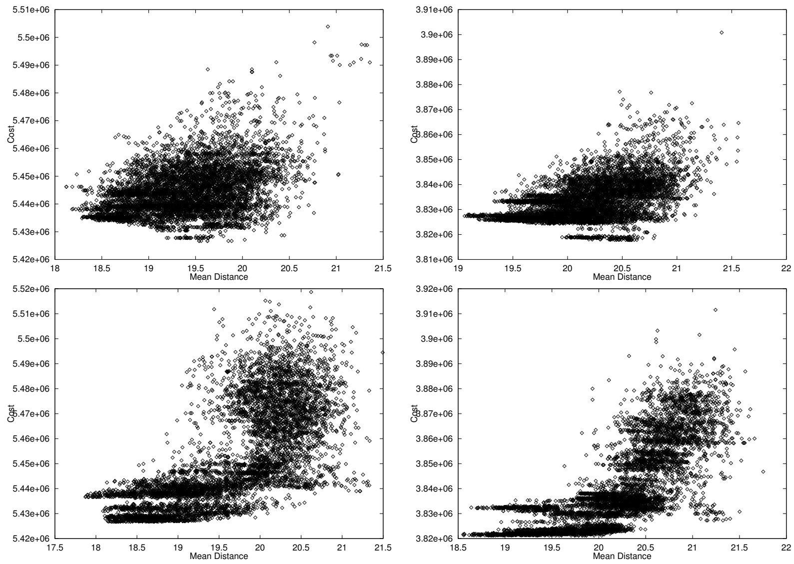

Fig. 2. No Big Valley structures of 6400 local optima from QAP instances. From left to right and top to down: bur26a, bur26b, bur26c, and bur26d.

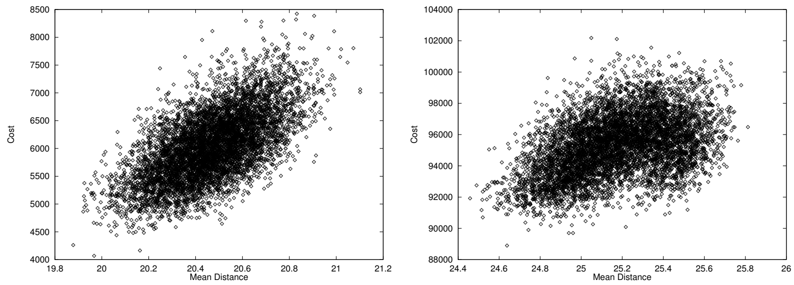

Fig. 3. Pseudo Big Valley structures of 6400 local optima from QAP instances. From left to right: chr25a and kra30a.

5.2 Perturbation vs. local optimum structures

For the experimental analysis of each instance, we used a random solution as the initial one, a greedy first improvement as local search, the location swapping of two facilities as neighborhood and two acceptance criteria: accepting only the current best local optimum and accepting any local optimum. In order to test the influence of local optimum structures on perturbation size, we use a simple random walk of different lengths as perturbation method in our ILS algorithm.

Table 2 shows the mean relative error ($(MILS - Best)/Best$), where $MILS$ is the mean value of 50 ILS runs, accepting only the current best local optimum,

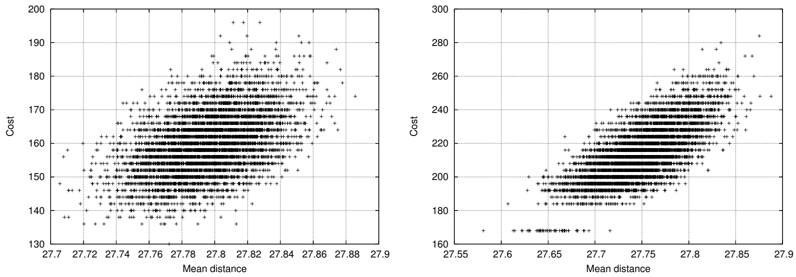

Fig. 4. Pseudo Big Valley structures of 6400 local optima from QAP instances. From left to right: esc32a and esc32b.

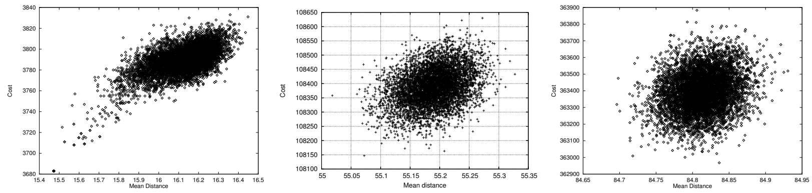

Fig. 5. Different structures of 6400 local optima from QAP instances. From left to right: lipa20a, lipa60a, and lipa90a.

and *Best* is the best known solution. These results are presented for different perturbation sizes from $N/16$ to N. We can see that for some instances we obtain a better result using a larger perturbation (*i.e.* close to N) while for others a smaller perturbation seems to be more appropriate (*i.e.* close to $N/16$). If we relate these results to the local optimum structures previously presented we can see that instances having better results with smaller perturbations have a clear BV structure, while instances that need larger perturbations have structures that differ from the BV. For each instance we take the best achieved result (the one in bold) and compare it with the value obtained with the farthest walk length to verify if the differences are statistically significant. For instance, for bur26a the best value was obtained with $N/2$, then we compare it with the value obtained for a walk length of $N/16$. For all instances the confidence intervals were calculated and they do not overlap (at a 99% of confidence), the exception was esc32b for which differences between averages results obtained for different walk lengths are not statistically significant.

Table 3 shows equivalent results when the accepting criterion is to accept any local optimum. Relation between BV structure and perturbation seems to hold for this criterion too, regardless of solutions' quality. Moreover, for instances having the BV structure and obtaining better results with smaller perturbations, we obtain better mean results accepting only the current best local optimum than accepting any local optimum (statistically tested at 99% of confidence).

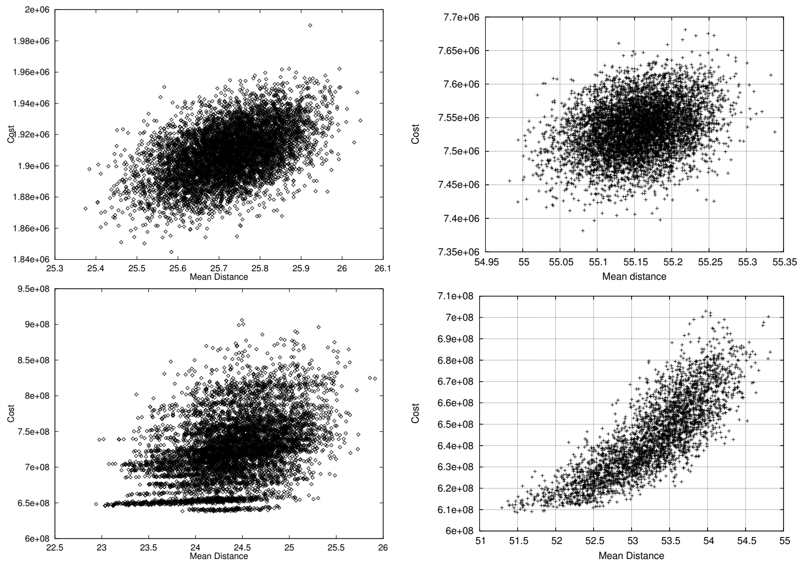

Fig. 6. Different structures of 6400 instance local optima. From left to right and top to down: tai30a, tai60a, tai30b, and tai60b.

Table 2. Relative errors between mean ILS results and best known solution for different values of walk lengths (accepting only the current best local optimum).

Name	**N/16**	**N/8**	**N/4**	**N/2**	**N**
bur26a	0.145	0.090	0.077	**0.067**	0.082
bur26b	0.182	0.135	0.123	0.105	**0.094**
bur26c	0.117	0.082	0.019	**0.003**	0.005
bur26d	0.109	0.093	0.027	**0.003**	0.005
chr25a	23.678	**19.589**	20.155	24.282	26.563
esc32a	6.2	**5.769**	7.585	8.954	9.631
esc32b	9.786	7.786	**7.167**	7.929	8.524
kra30a	3.678	2.632	**1.901**	2.332	2.623
lipa20a	1.916	1.491	1.325	**1.122**	1.312
lipa60a	**0.828**	0.858	0.897	0.932	0.941
lipa90a	**0.587**	0.622	0.663	0.677	0.679
tai30a	2.575	2.204	**2.059**	2.575	2.813
tai30b	4.150	2.373	1.388	**0.759**	0.797
tai60a	2.407	**2.190**	2.980	3.132	3.063
tai60b	1.626	1.004	0.734	**0.376**	0.412

This confirms the existence of BV structure, since accepting only the current best local optimum leads to a biased search that exploits this structure and allows better results using smaller perturbations.

Table 3. Relative errors between mean ILS results and the best known solution for different values of walk lengths (accepting any local optimum).

Name	**N/16**	**N/8**	**N/4**	**N/2**	**N**
bur26a	0.148	0.126	0.101	**0.078**	0.080
bur26b	0.182	0.142	0.109	**0.080**	0.092
bur26c	0.168	0.036	0.007	**0.005**	0.008
bur26d	0.122	0.048	**0.005**	0.009	0.006
chr25a	29.299	**24.228**	27.118	25.555	26.769
esc32a	**7.538**	7.846	8.246	9.261	9.385
esc32b	8.905	7.286	**6.143**	7.048	7.619
kra30a	3.537	2.530	**2.162**	2.339	2.747
lipa20a	1.632	1.386	**1.155**	1.527	1.749
lipa60a	0.895	**0.879**	0.921	0.933	0.943
lipa90a	**0.633**	0.661	0.668	0.672	0.676
tai30a	2.474	**2.307**	2.710	2.752	2.889
tai30b	6.862	2.510	1.568	0.973	**0.755**
tai60a	**2.495**	2.884	3.009	3.171	3.167
tai60b	0.723	0.950	**0.179**	0.414	0.671

There are some instances for which the relation between local optimum structures and perturbation length is not clear. For instance, tai30a in Figure 6 seems to have a better BV structure than lipa90a in Figure 5, contradicting the results shown in Table 2, that indicates that lipa90a needs a smaller perturbation. This could be a visualization problem which does not allow us to conclude only by considering the cost-distance graphs. We need a measure that can tell us what is happening in these cases.

If we want to measure how much an instance have a BV structure, we need to take into account two things: cost-distance correlation and closeness of local optima. Table 4 shows two measures for each instance. The first one is the Pearson correlation [7] between cost and distance of the 6400 local optima obtained in Section 5.1. Comparing these values with ILS results we can see that there is no relation between them, since there are instances that have almost the same correlation but need different perturbations lengths (*i.e.* bur26b and esc32a). The second measure is obtained by taking the difference between the maximum and minimum mean distance (DMD) divided by problem size (N). Relating this measure with Table 2 we can see that as the ratio DMD/N decreases, a smaller perturbation is needed, and viceversa, as relation DMD/N increases, a larger perturbation is needed. That is, DMD/N values have a direct relation to perturbation length, even for instances for which we can not conclude the same by visualizing only cost-distance graphs. Such relation seems to be $(N * (DMD/N)) * N = DMD * N$ as an upper bound in the random walk length needed as perturbation in ILS to obtain good results.

Table 4. Cost-Distance correlation and relative closeness of local optima.

Name	Correlation	DMD/N
bur26a	0.387	0.135
bur26b	0.494	0.115
bur26c	0.702	0.154
bur26d	0.753	0.135
chr25a	0.625	0.056
esc32a	0.472	0.006
esc32b	0.723	0.011
kra30a	0.430	0.053
lipa20a	0.677	0.055
lipa60a	0.337	0.006
lipa90a	0.214	0.003
tai30a	0.485	0.027
tai30b	0.420	0.117
tai60a	0.293	0.007
tai60b	0.818	0.067

6 Conclusions

An analysis of local optimum structures for QAP instances has been presented. The analysis was focused on searching for a Big Valley structure in the set of local optima. The analysis shows that QAP instances have different cost-distance relations between their local optima, including BV. An experimental analysis was performed in order to relate these structures to the required perturbation length of the random walk. Experimental results show that instances that have BV structure need an smaller perturbation to reach better values, while instances having structures far from a BV need a larger perturbation. In the case of those instances for which a clear relation between local optimum structures and perturbation values could not be established, a relation between local optimum closeness and problem size has been introduced to explain the phenomenon successfully. These results will help in the design of ILS based algorithms.

Future research is planned to extend these results to a wider set of instances, and to use different local searches such as Simulated Annealing, Tabu Search, and others, as part of the ILS algorithm. Another obvious future work is to verify whether the relation between BV and perturbation holds for other combinatorial optimization problems like the TSP. Finally, a deeper analysis is needed to exploit each local optima structure in order to design a specific combination of ILS mechanisms to obtain better results.

Acknowledgments

The authors would like to thank the anonymous referees for their useful comments and interesting ideas for future research.

References

1. E. H. L. Aarts and J. K. Lenstra, editors. *Local Search in Combinatorial Optimization*. Wiley, Chichester, 1997.
2. K. D. Boese. *Models for Iterative Global Optimization*. PhD thesis, University of California at Los Angeles, Los Angeles, CA, 1996.
3. K. D. Boese, A. B. Kahng, and S. Muddu. A new adaptive multi-start technique for combinatorial global optimization. *Operations Research Letters*, 16:101–113, 1994.
4. T. Dimitriou and R. Impagliazzo. Towards a rigorous analysis of local optimization algorithms. In *25th ACM Symposium on the Theory of Computing*, 1996.
5. D. Du and Panos M. Pardalos, editors. *Handbook of Combinatorial Optimization*. Kluwer Academic Publishers, London, 1999.
6. M. R. Garey and D. S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-completeness*. W. H. Freeman, San Francisco, 1979.
7. Paul G. Hoel, editor. *Introduction to Mathematical Statistics*. John Wiley & Sons, New York, 1962.
8. T. C. Hu, V. Klee, and D. Larman. Optimization of globally convex functions. *SIAM Journal on Control and Optimization*, 27(5):1026–1047, 1989.
9. D. S. Johnson and L. A. McGeoch. The travelling salesman problem: A case study in local optimization. In E. H. L. Aarts and J. K. Lenstra, editors, *Local Search in Combinatorial Optimization*, pages 215–310. Wiley, 1997.
10. T. Jones. *Evolutionary Algorithms, Fitness Landscape and Search*. PhD thesis, The University of New Mexico, Albuquerque, New Mexico, 1995.
11. H. R. Lourenco, O.C. Martin, and T. Stutzle. Iterated local search. In F. Glover and G. Kochenberger, editors, *Handbook of Metaheuristics*. Kluwer, 2002.
12. B. Manderick, M. De Weger, and P. Spiessens. The genetic algorithm and the structure of the fitness landscape. In R. K. Belew, editor, *Fourth International Conference on Genetic Algorithms*, pages 143–150, San Mateo, CA, 1991. Morgan Kaufmann.
13. O. Martin, S.W. Otto, and E.W. Felten. Large-step markov chains for the traveling salesman problem. *Complex Systems*, 5(3):299–326, 1991.
14. D. C. Mattfeld and C. Bierwirth. A search space analysis of the job shop scheduling problem. *Annals of Operational Research*, 86:441–453, 1999.
15. C. R. Reeves. Landscapes, operators and heuristic search. *Annals of Operational Research*, 86:473–490, 1999.
16. C. M. Reidys and P. F. Stadler. Combinatorial landscapes. *SIAM Review*, 44(1):3–54, 2002.
17. S. Russell and P. Norvig. *Artificial Intelligence: A Modern Approach*. Prentice Hall, 1995.
18. E. D. Taillard. Comparison of iterative searches for the quadratic assignment problem. *Location Science*, 3:87–105, 1995.
19. E. Weinberger. Correlated and uncorrelated fitness landscapes and how to tell the difference. *Biological Cybernetics*, 63:325–336, 1990.