

Towards Landscape Analyses to Inform the Design of a Hybrid Local Search for the Multiobjective Quadratic Assignment Problem

Joshua D. Knowles and David W. Corne

Abstract. The quadratic assignment problem (QAP) is a very difficult and practically relevant combinatorial optimization problem which has attracted much research effort. Local search (LS) moves can be quickly evaluated on the QAP, and hence favoured methods tend to be hybrids of global optimization schemes and LS. Here we introduce the *multiobjective* QAP (mQAP) where $m \geq 2$ distinct QAPs must be minimized simultaneously over the same permutation space, and hence we require a set of solutions approximating the Pareto front (PF). We argue that the best way to organise a hybrid LS for the mQAP will depend on details of the multiobjective fitness landscape. By using various techniques and measures to probe the landscapes of mQAPs, we attempt to find evidence for the relative ease with which the following can be done by LS: approach the PF from a random initial solution, or search along or close to the PF itself. On the basis of such explorations, we hope to design an appropriate hybrid LS for this problem. The paper contributes a number of landscape measurement methods that we believe are generally appropriate for multiobjective combinatorial optimization.

1 Introduction

Many problems can be formulated as a *quadratic assignment problem* (QAP). Computationally, this is NP hard [10] and even relatively small general instances ($n \geq 20$) cannot be solved to optimality. We introduce here the *multiobjective* QAP (mQAP) as both a practically important problem and a benchmark that we believe will be useful in general research on multiobjective combinatorial landscapes. As in the scalar QAP, fast local search (LS) is possible because a difference in the objective function value(s) can be calculated quickly when using a 2-exchange operator.

The question of interest is how to control and organise the many required runs of LS. Depending on the landscape, it may not be efficient to start each search *ab initio*. Rather, it may be better to search in parallel for the different Pareto optima, by alternating the ‘search direction’ in the objective space, step by step. Or, alternatively again, the best approach may be to first find one Pareto optimum and then to search along the Pareto front from this one to find others.

Our aim here is to propose some multiobjective landscape analysis methods that might be used to predict the best approach given an unseen mQAP instance. The paper is necessarily speculative as little work has been done on the best way to apply hybrid LS in multiobjective domains and still less on multiobjective landscape analysis.

The remainder is organized as follows. Section 2 reviews the scalar QAP and describes a best improvement LS that can also be used for the mQAP. Section 3 introduces the mQAP, proposing applications and its theoretical interest. A test instance generator is described in Section 4. Section 5 outlines different overall search strategies for approximating the PF of the mQAP, relating these to landscape correlations. Tools for measuring properties of the multiobjective landscape are described in Section 6, and some illustrative results are given. Section 7 concludes.

2 The scalar QAP

The quadratic assignment problem (QAP) is the problem of assigning each of n facilities to exactly one of n possible locations so as to minimize the sum of products of flows between the facilities and the distances between the locations. For a mathematical formulation of the problem see that for the mQAP in section 3, the only difference being the use of multiple ($m \geq 2$) flow matrices. Solutions to the QAP may be represented by a permutation π on the set $\{1, \dots, n\}$, where n is the number of facilities/locations.

A deterministic best-improvement local search for the QAP, based on exchanging the locations of two facilities, was proposed by Taillard [3]. The search is fast because it is possible to calculate the change in cost of an exchange in $O(n)$ time as only the two affected rows and columns in the flow matrix must be re-evaluated. Further, it is possible to use solutions from previous iterations to calculate some of the exchanges in just $O(1)$ time.

Much work has been concerned with measures on the fitness landscape of QAP instances. Here we will just give those used later in this paper, which mainly depend on a measure of distance in the parameter (permutation) space. For example, Bachelet [1] measures distance $\text{dist}(\pi, \mu)$ between two solutions π and μ as the smallest number of 2-swaps that must be performed to transform one solution into the other (this can be computed in $O(n)$ time); this distance measure has a range of $[0, n - 1]$. From this, other measures can be easily defined. Bachelet gives the diameter of a population P as

$$\text{dmm}(P) = \frac{\sum_{\pi \in P} \sum_{\mu \in P} \text{dist}(\pi, \mu)}{|P|^2}.$$

The entropy [4], which is a further measure of the dispersion of solutions, is given by

$$\text{ent}(P) = \frac{-1}{n \log n} \sum_{i=1}^n \sum_{j=1}^n \left(\frac{n_{ij}}{|P|} \log \frac{n_{ij}}{|P|} \right)$$

where n_{ij} is the number of times facility i is assigned to location j in the population. The fitness distance correlation [9] has been widely used. Often a plot is shown, either giving the fitness against distance from the nearest global optimum or best known solution, or all pairs of solutions may also be differenced.

Vollmann and Buffa [12] introduced flow dominance as a means of characterizing QAP instances. Flow dominance is a measure of the flow matrix given by

$$fd(c) = 100 \times \frac{a}{b} \quad a = \sqrt{\frac{\sum_{i=1}^n \sum_{j=1}^n (c_{ij} - b)^2}{n^2}} \quad b = \frac{\sum_{i=1}^n \sum_{j=1}^n c_{ij}}{n^2}$$

When there is high flow dominance there is low epistasis, in general. The distance dominance dd can be defined on the distance matrix in an analogous fashion.

3 The multiobjective QAP model

The multiobjective QAP (mQAP), with multiple flow matrices (defined below) naturally models any facility layout problem where we are concerned with the flow of more than one type of item or agent. For example, in a hospital layout problem we may be concerned with simultaneously minimizing the flows of doctors on their rounds, of patients, of hospital visitors, and of pharmaceuticals and other products. We may imagine that each of the matrices describing these flows is very different and that until we see the Pareto front it may not be possible to simply weight the importance of each flow. Thus we have a very natural multiobjective problem. Similar examples can be given for other facility location problems such as the layout of factory floors and the topology of distribution networks. Furthermore, other problems that can be modeled using QAP, such as scheduling, may also exist in multiobjective form. The mQAP is also likely to be a useful problem to study from a theoretical viewpoint.

Formally, the mQAP problem is to ‘minimize’:

$$\vec{c}(\pi) = \{c^1(\pi), c^2(\pi), \dots, c^m(\pi)\}, \text{ where } c^k(\pi) = \sum_{i=1}^n \sum_{j=1}^n a_{ij} b_{\pi_i \pi_j}^k, \quad k \in 1..m$$

where n is the number of facilities/locations, m is the number of different flow matrices, a_{ij} is the distance between location i and location j , b_{ij}^k is the k th flow from facility i to facility j , and π_i gives the location of facility i in permutation $\pi \in P(n)$, where $P(n)$ is the set of all permutations of $\{1, 2, \dots, n\}$.

The term ‘minimize’ must also be defined because different methods are available for tackling multiobjective problems, including lexicographic ordering, scalarizing techniques, or constraint methods. However, it is often undesirable to use any of the aforementioned techniques, which induce a total ordering on the solutions, without knowledge of what trade-offs of the objectives are ‘out there’ to be found. In this case it is best to use Pareto optimization, a technique that makes the fewest possible assumptions about which solutions are preferable, before the search is conducted. Here,

the aim of Pareto optimization is to find a set of solutions $P^*(n) \subseteq P(n)$, the Pareto optimal set, which are *nondominated* in the whole search space:

$$\forall \pi^* \in P^*(n) \nexists \pi \in P(n) \text{ such that } \pi \prec \pi^*, \text{ where } \pi \prec \pi^* \text{ iff} \\ \forall k \in \{1, \dots, m\} c^k(\pi) \leq c^k(\pi^*) \wedge \exists k \in \{1, \dots, m\} \text{ such that } c^k(\pi) < c^k(\pi^*)$$

where $\pi \prec \pi^*$ is read as π *dominates* π^* . The set of objective vectors (or points) $C^* \subseteq C$, called the Pareto front (PF), is the image of the Pareto optimal set in the objective space, $\vec{c}(P^*(n)) \mapsto C^*$, and C is the image of the whole permutation space $\vec{c}(P(n)) \mapsto C$. The aim of Pareto optimization is to find a good approximation to the whole Pareto front C^* . To understand what is meant by a good approximation, the reader is referred to [7] but need only know that sets of solutions (approximation sets) can be partially ordered without additional assumptions.

4 An mQAP instance generator

To obtain instances for the mQAP we have developed a generator that allows different problem features to be controlled. The generator makes symmetric multi-objective QAP instances with one distance matrix and multiple flow matrices. The distance and flow matrices have structured, ‘real-world-like’ entries [3] and correlations can be set between corresponding entries in the two or more flow matrices. The generator follows procedures for making the non-uniformly random QAP problems given the appellation TaiXXb in the literature, outlined in [3].

The number of flow matrices is controlled by a parameter, m . With $m \geq 2$, a multiobjective QAP problem is generated with m flow matrices. The entries in the k th matrix ($2 \leq k \leq m$) are generated where a random variable X is correlated with the value of X used in the corresponding entry in the first flow matrix. Correlations between -1 and 1 can be set between the first and each of the additional flow matrices using $m - 1$ further parameters to the generator.

A degree of overlap between the matrices can also be specified. The overlap parameter is set between 0 and 1 . It controls the fraction of entries in the j th flow matrix that are correlated with the corresponding entries in the 1st flow matrix. With the overlap parameter set to zero, a random un-correlated value will be placed in each entry of the j th flow matrix that corresponds to a zero entry in the first flow matrix. Similarly, a zero will be placed in each entry of the j th flow matrix that corresponds to a non-zero value in the first flow matrix. Thus there is no overlap between the flows of the first and j th matrix. With the overlap parameter set to 1 all the flows are correlated. With the overlap set to intermediate values some of the flows will overlap and others will not.

5 Hybrid search schemes for the mQAP

Multiobjective (MO) metaheuristics for Pareto optimization have been around for quite some time now and MO versions of genetic algorithms [5], simulated annealing [2],

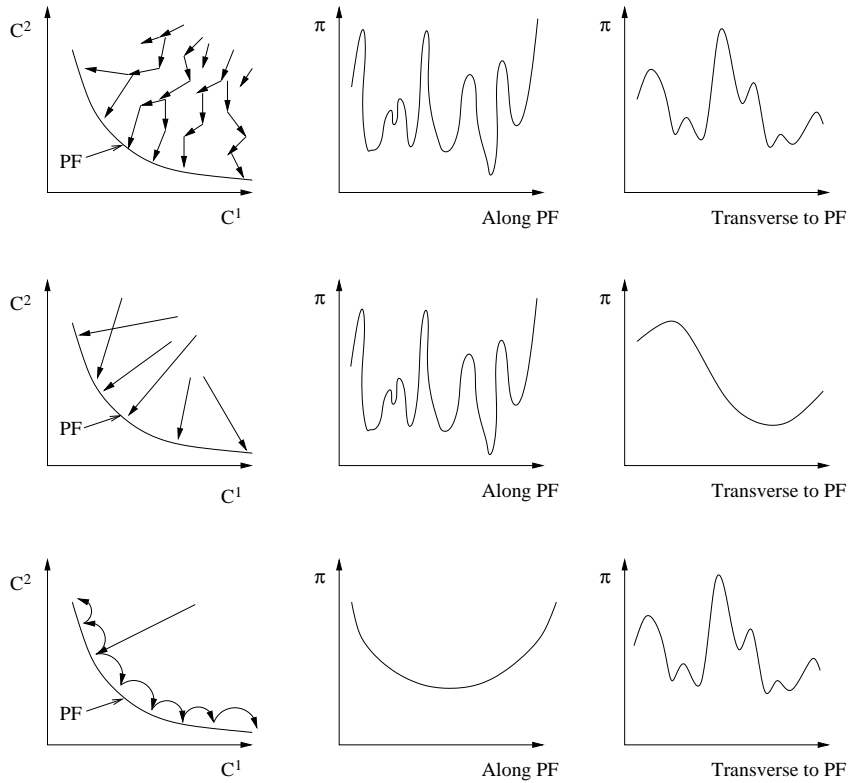


Figure 1: Ways to approach the Pareto front and the underlying landscape appropriate to each. The left column shows the Pareto front of a two-objective problem and different ways a search algorithm may find points close to the PF. The centre and right columns depict correlations between the multiobjective fitness and the distance between points in the parameter space. See main text for further explanation

tabu search (TS) [6], and others have been proposed. Some hybrid genetic algorithm approaches have also been put forward, e.g. [8]. With the mQAP, the availability of a fast best-improvement LS (or TS) means that any hybrid approach should consider employing LS by specifying a scalarizing weight vector and allowing the LS to find a local optimum of this search ‘direction’. The question is, how should the runs of the LS be controlled by the overall search strategy? This question is already difficult in scalar optimization, We believe the answer concerns the best way to move about in the multiobjective landscape.

Figure 1 illustrates three quite different strategies, using a 2-objective problem as an example. The first strategy is to use a population-based approach which tries to gradually improve the population in all directions at once, approaching the PF from all angles, with LS applied to all solutions in the population at each generation. In the second strategy, a scalarizing weight vector is first chosen and the LS is applied multiple times in only this direction until no more improvement results. Each LS restart could be from a random point but more usually it would be from a perturbation

of the previous local optimum, taking it out of that basin of attraction [11]. After one point on or near the PF is found the process is repeated for a different scalarizing vector. In the last strategy we envisage moving towards the PF in one direction by repeated LS applications as in the previous method. But then once on the PF, the same solution should be perturbed again and minimized over a ‘nearby’ scalarizing vector, resulting in a point nearby on the PF.

Clearly, the best approach will depend on whether or not points nearby in objective space are also nearby in the permutation space. In particular, we can consider this in two different ways. For a particular scalarizing vector is it easy to move towards (perpendicular to) the PF i.e. is the fitness distance correlation good in this search direction? And for a set of mutually nondominated points (i.e for points sharing a common ‘fitness’ but not in the same objectives), is it easy to move from one point to others nearby, parallel to the local Pareto front? It is possible that the degree of correlation in the landscape varies tremendously depending on which direction of search is considered and/or what is the fitness level. It may be very difficult to find meaningful measurements of these properties and even more difficult to predict which will be the best method of search. Nonetheless, it seems likely that the success of searches will depend heavily on what kind of overall strategy is employed, and this in turn on the landscape, so it seems well worth trying to gain some understanding of local correlations as well as global properties. In the next section we describe some example measurements we have made and discuss what they could mean.

6 Landscape measurement tools and example results

Our methods are based on those previously proposed for the scalar QAP, as briefly reviewed in section 2.1 (and thoroughly described in [1]). The overall strategy in the scalar QAP case is to run a local search from several random starting solutions. Measuring different properties of the local optima obtained, as well as properties of the starting solutions, gives some insight into how correlated the optima are, and thereby hints at the difficulty of finding a global optimum. The same strategy is applicable to the mQAP, although complicated by the fact we need to try to find optima all over the PF.

Algorithm 1 shows our approach. We start by defining a set of evenly distributed scalarizing vectors Λ . For each of these ‘directions’ we repeat several LS runs from random start solutions. For each such LS run, we record the starting point, its multiobjective evaluation, the local optimum reached, its multiobjective evaluation, the scalarizing weight vector used in the LS, and the number of LS moves applied. We then analyse these records. First, we can just compare some properties of all starting solutions, U , and optima O : the entropy, the diameter and the number of LS moves performed on U to obtain O , $\text{lmm}(U) = \text{lmm}(O)$. Note that although entropy and diameter were defined for the QAP [1], these can be used without alteration on the mQAP. Next, we can try to investigate properties of the landscape perpendicular and

Table 1: Results of our measurements on two 20 node, 2 flow matrix, and two 30 node, 3 flow matrix instances with different flow correlations and other parameters. The first figure in an instance name is the number of nodes, the second is the number of flow matrices and the third is just an index. The global properties are: $\max(d)$, the maximum distance in the distance matrix; $\max(f)$ the maximum flow in any of the flow matrices; $\text{corr}(f^i, f^j)$ the correlation between corresponding flow matrix entries of the i th and j th flow; the overlap of correlated entries; fd^k , the flow dominance of the k th flow matrix; and dd the distance dominance. The local search measures were collected from 100000 (105000, 3 objective) starting points. Figures in brackets are sample sizes. Where U or O is primed this means that the measure is calculated on a sample of this set. The sampling criteria are as follows: rs refers to a random sample; $\vec{\lambda}$ to solutions satisfying the condition, $\sum_{j=1}^k (\lambda^k - \xi^k)^2 < 0.01$ where $\xi = (0.0, 0.3, 0.7)$ for the three objective instances, and $\xi = (0.1, 0.9)$ for the two-objective instances; and nds refers to the nondominated solutions from the full set

Glob. Property	Instances			
	Kno20-2fl-1rl	Kno20-2fl-2rl	Kno30-3fl-1rl	Kno30-3fl-2rl
$\max(d)$	196	173	172	180
$\max(f)$	9883	9644	9967	9968
$\text{corr}(f^1, f^2)$	0.0	0.4	0.4	0.7
$\text{corr}(f^1, f^3)$	–	–	0.0	-0.5
overlap	1.0	1.0	0.7	0.7
fd^1, fd^2, fd^3	206,260,–	243,315,–	235,354,252	234,324,342
dd	54.9	57.9	56.1	58.6
LS Measure				
$\text{ent}(U)$	0.999971	0.999971	0.999959	0.999962
$\text{ent}(O)$	0.910517	0.870681	0.938710	0.915543
$\text{diam}(U'_1, rs)$	16 (10^3)	16 (10^3)	23 (10^3)	24 (10^3)
$\text{diam}(O'_1, rs)$	16 (10^3)	16 (10^3)	23 (10^3)	23 (10^3)
$\text{lmm}(U'_1, rs)$	14.73 (10^3)	15.459 (10^3)	26.007 (10^3)	26.239 (10^3)
$\text{diam}(O'_1, \vec{\lambda})$	16 (10^3)	16 (10^3)	24 (10^3)	23 (10^3)
$\text{ent}(O'_1, \vec{\lambda})$	0.900 (10^3)	0.847 (10^3)	0.853 (10^3)	0.846 (10^3)
$\text{lmm}(O'_2, \vec{\lambda})$	13.719 (10^3)	15.511(10^3)	25.625(10^3)	26.866 (10^3)
$\text{diam}(O'_3, nds)$	14 (200)	14 (200)	22 (200)	22 (200)
$\text{ent}(O'_3, nds)$	0.489 (200)	0.252 (200)	0.743 (200)	0.746 (200)
$\text{lmm}(O'_3, nds)$	16.695 (200)	17.310 (200)	28.945 (200)	29.955 (200)

parallel to the PF. To do this we construct samples (subsets) of a complete set of records (i.e. of O or U or $O \cup U$) such that all records in the sample satisfy some condition on them. Three different conditions, and combinations of these, seem promising. The first is to consider only optima that have been found using a similar LS direction, which we can specify using a target vector ξ and the following condition: $\sum_{j=1}^k (\lambda^k - \xi^k)^2 < \epsilon$, where ϵ is some small positive value. The second condition is to take only solutions of the same nondominated rank. In the simplest case this just means taking only the nondominated solutions (i.e., Pareto optimal from among those points found). The

Algorithm 1 Outline of a landscape analysis tool for the mQAP

```
1: Input:  $n$ ;  $\Lambda$ ;  $\#searches$ ; mQAP instance  $I$ ; conditions  $C$ ; measures  $M$ 
2:  $U \leftarrow \emptyset$ 
3:  $O \leftarrow \emptyset$ 
4: for each  $\vec{\lambda} \in \Lambda$  do
5:   for  $i = 1$  to  $\#searches$  do
6:      $\pi \leftarrow \text{RandomPermutation}(n)$ 
7:      $\pi' \leftarrow \text{LS}(\pi, I)$ 
8:      $r(\pi) \leftarrow (\pi, \vec{C}(\pi), \#LSmoves, \vec{\lambda})$ 
9:      $r(\pi') \leftarrow (\pi', \vec{C}(\pi'), \#LSmoves, \vec{\lambda})$ 
10:     $U \leftarrow U \cup \{r(\pi)\}$ 
11:     $O \leftarrow O \cup \{r(\pi')\}$ 
12:   end for
13: end for
14: for each measure  $mre \in M$  do
15:   print  $mre(U)$ 
16:   print  $mre(O)$ 
17: end for
18: for each condition  $c \in C$  do
19:    $S \leftarrow \{r(\pi) \in T \mid c(r(\pi)) = \text{TRUE}\}$ , where  $T = O \vee U \vee O \cup U$ 
20:   for each measure  $mre \in M$  do
21:     print  $mre(S)$ 
22:   end for
23: end for
```

third condition is to specify some target objective vector and take only points nearby to this. By measuring the entropy and diameter of these subsets we can get a feeling for how similar solutions within these subsets are. We can of course combine these partitioning conditions together to further our investigations. We are also able to measure fitness-distance correlations within our different subsets using, as fitness, the Manhattan distance between points in the objective space.

Table 1 presents some illustrative results of our landscape analyses for five different mQAP instances. These were collected by running 1000 local searches from each of 100 (for 2-objective instances) or 105 (3-objective instances) different λ vectors, thus giving us approximately 200000 records in all. Although at this stage it is not possible to draw any conclusions about which overall search strategy would be most suitable for each instance, we can see that pertinent information is provided. For example, we can see that the diameter of the subsets tends to be about the same size as that of all of the entire sets U and O . Entropy, however, changes, indicating that although optima can be found over a large region, they are at least clustered together in groups. For the 2-objective instances, the entropy of the nondominated points is quite low. This means, particularly in the instance with positive correlation (0.4) in the flow matrix entries,

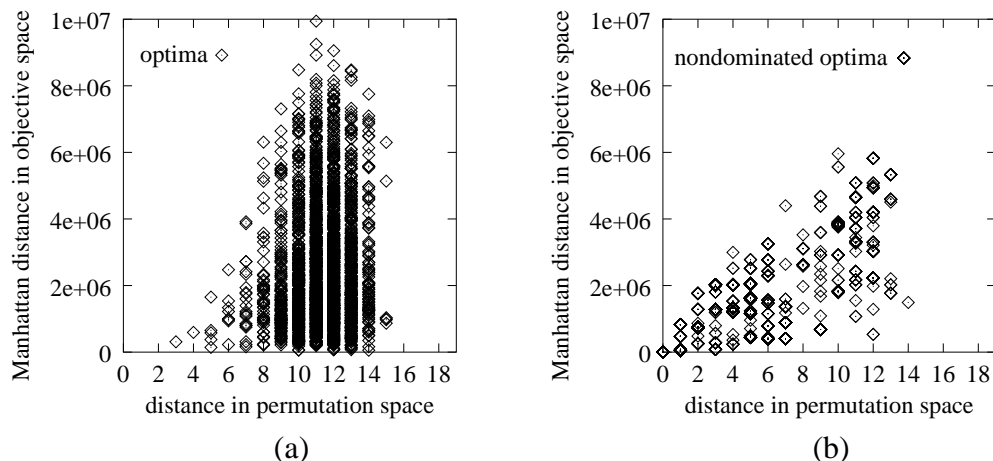


Figure 2: Fitness-distance correlation plots between all pairs of points. In (a) a random sample of local optima is plotted. In (b) a sample (of the same size) of nondominated optima is plotted. The greater fitness-distance correlation in (b) is clear

that points on the nondominated front are quite similar. This suggests that searching parallel to the PF may be much more efficient than re-starting searches from random points. Points found using a particular λ vector, or one close to it, do not have such a low entropy. This might suggest that on the way towards the PF optima are spread all over the space and are only in relatively small clusters. Thus from the preliminary results seen here we might guess that it is easier to move along the PF than to get to the PF in the first place. Many more experiments using these techniques are needed to confirm this hypothesis and to make more detailed inferences or predictions.

In Figure 2 we can see two fitness-distance correlation plots for (a) all pairs of a random sample of optima and (b) all pairs of the nondominated optima of instance Kno20-2fl-2rl. Clearly, the nondominated optima are highly correlated, indicating that it is probably quite easy to find solutions on the PF once one has been found.

7 Conclusion

Some problems with multiple objectives can be searched using a hybrid approach which makes use of an efficient local search. When deciding on an overall search strategy, however, we argue that one must consider whether there is correlation between nearby optima or not. In a multiobjective landscape we can measure such correlations either ‘perpendicular’ or ‘parallel’ to the Pareto front. This will help us decide whether our search should first approach the PF and then spread around from there, or whether our search should start repeatedly from new random points, or whether it should use a gradual approach towards the PF from all directions in parallel. In this paper we have introduced and discussed these notions and given some methods for measuring the landscape of a multiobjective combinatorial optimization problem, mQAP, which

we also introduced here for the first time.

Acknowledgments

Joshua Knowles gratefully acknowledges the support of a Marie Curie research fellowship of the European Commission, contract number HPMF-CT-2000-00992.

References

- [1] Vincent Bachelet. *Métaheuristiques Parallèles Hybrides: Application au Problème D'affectation Quadratique*. PhD thesis, Université des Sciences et Technologies de Lille, December 1999.
- [2] P. Czyzak and A. Jaskiewicz. Pareto Simulated Annealing. In G. Fandel and T. Gal, editors, *Multiple Criteria Decision Making. Proceedings of the XIIth International Conference*, pages 297–307, Hagen, Germany, 1997. Springer-Verlag.
- [3] Éric D Taillard. A comparison of iterative searches for the quadratic assignment problem. *Location Science*, 3:87–105, 1995.
- [4] C. Fleurent et J.A. Ferland. Genetic hybrids for the quadratic assignment problem. *DIMACS Series in Discrete Mathematics and Theoretical Computer Science*, 16:173–188, 1994.
- [5] Carlos M. Fonseca and Peter J. Fleming. Genetic algorithms for multiobjective optimization: Formulation, discussion and generalization. In Stephanie Forrest, editor, *Proceedings of the Fifth International Conference on Genetic Algorithms*, pages 416–423, San Mateo, California, 1993. University of Illinois at Urbana-Champaign, Morgan Kaufmann Publishers.
- [6] Michael Pilegaard Hansen. Tabu search in multiobjective optimisation : MOTS. In *Proceedings of the 13th International Conference on Multiple Criteria Decision Making (MCDM'97)*, Cape Town, South Africa, January 1997.
- [7] Michael Pilegaard Hansen and Andrzej Jaskiewicz. Evaluating the quality of approximations to the non-dominated set. Technical Report IMM-REP-1998-7, Institute of Computing Science, Poznan University of Technology, ul. Piotrow 3a, 60-965 Poznan, Poland, March 1998.
- [8] Andrzej Jaskiewicz. On the performance of multiple objective genetic local search on the 0/1 knapsack problem. a comparative experiment. Technical Report RA-002/2000, Institute of Computing Science, Poznan University of Technology, Poznań, Poland, July 2000.
- [9] T. Jones and S. Forrest. Fitness distance correlation as a measure of problem difficulty for genetic algorithms. In L. J. Eshelman, editor, *Proceedings of the 6th International Conference on Genetic Algorithms*, pages 184–192. Morgan Kaufmann, 1995.
- [10] S. Sahni and T. Gonzalez. P-complete approximation problems. *Journal of the ACM*, 23:555–565, 1976.
- [11] Thomas Stützle. Iterated local search for the quadratic assignment problem. Technical Report AIDA-99-03, Darmstadt University of Technology, Computer Science Department, Intellectics Group, 1999.
- [12] T. E. Vollmann and E. S. Buffa. The facilities layout problem in perspective. *Management Science*, 12(10):450–468, 1966.