A Hospital Facility Layout Problem
Finally Solved

by
Peter M. Hahn
Department of Systems Engineering
University of Pennsylvania
293 Towne Building, 220 S. 33rd Street
Philadelphia, PA 19104-6315, USA
E-mail: peter_hahn@compuserve.com
and
Jakob Krarup
DIKU, Department of Computer Science
University of Copenhagen
Universitetsparken 1, DK-2100
Copenhagen, Denmark
E-mail: krarup@diku.dk

April, 2000

ABSTRACT

This paper presents a history of a difficult facility layout problem that falls into the category of the Koopmans-Beckmann variant of the Quadratic Assignment Problem (QAP), wherein 30 facilities are to be assigned to 30 locations. The problem arose in 1972 as part of the design of a German university hospital, Klinikum Regensburg. This problem, known as the Krarup30a upon its inclusion in the QAPLIB library of QAP instances, has remained an important example of one of the most difficult to solve. In 1999, two approaches provided multiple optimum solutions. The first was Thomas Stützle's analysis of Fitness-Distance Correlation that resulted in the discovery of 256 global optima. The second was a new branch-and-bound enumeration that confirmed 133 of the 256 global optima found and proved that Stützle's 256 solutions were indeed optimum. We report here on the steps taken to provide in-time heuristic solutions and the methods used to finally prove the optimum.

KEYWORDS: assignment problems, combinatorial optimization, facility location, facility planning, hospital planning, integer programming, quadratic assignment problem.

1. The Problem

Spadille Inc., Consultants of Industrial Statistics and Operations Research was established in 1970. Jakob Krarup was one of the three co-founders and served as a manager of this still existing Danish company until his return to academia in 1975.

Among the projects undertaken was the design of a university hospital, "Klinikum Regensburg" to be built in Regensburg, Germany. An invitation to submit tenders was issued to a number of architects. As was explicitly stated in the announcement of the competition, among the criteria to be considered in the evaluation of proposals was the usual QAP objective: find a layout minimizing the sum CDIST of Communication x DISTance taken over all pairs of facilities to be located.

To assess the architects' proposals regarding this criterion only, Spadille's task was to find "a best lower bound" on CDIST. Or, if at all possible, Spadille was to find provably optimal solutions to a series of instances of the Koopmans-Beckmann variant of the more general QAP. Included were two instances which later were to be placed in QAPLIB, a Quadratic Assignment Problem Library, established in 1991 by Burkard, Karisch and Rendl (Burkard, 1991) as the "Krarup 30a/b".

Let n be the number of *facilities* to be located in n out of m *cells,* m>n. For a hospital, facilities can be viewed as equally sized objects or "functions" like surgery, X-ray, et cetera. For our purpose, we can consider a *cell* as any point with positive integer coordinates in a 3-dimensional coordinate system. Each cell can accommodate exactly one of the n facilities to be located.

It can be argued that the two cases "m=n" and "m>n" essentially are equivalent since m-n "dummy" facilities can be added if m is strictly greater than n. From a practitioners' point of view, however, these two cases may well represent different situations.

With the straightforward interpretation of QAP in mind (assignment of facilities to cells within a building), m=n means that the building is designed to exactly house the n facilities in its free cells. Of course, there may also be other cells in the building, but they are in this context regarded as inadmissible, perhaps because they are reserved for other purposes.

The "m>n"-case is relevant when none or only some information is available as to the exact shape of the building. For example, for "Klinikum Regensburg", the only restrictions imposed were due to the surrounding landscape. Thus, a beautiful pond in the park had to be protected. Other aesthetic criteria to be taken into account were reflected by an upper bound on the number of floors. The shape of the building was not otherwise determined a priori but supposed to result from the "most compact arrangement of the n facilities" suggested by the architects or determined by solving the corresponding QAP.

2. Spadille's Approach to Solving the Problem

A number of instances with n ranging from 30 to 48 were investigated. Since no algorithm at that time (1972) was capable of solving such sizeable instances to optimality, Spadille could hope for nothing better than generating lower bounds, hopefully below the objective function values reached by the architects.

To this end, a *randomized heuristic* was devised (Krarup, 1972 and Krarup, 1978). Let x_u,y_u,z_u be upper bounds on the length, the width, and the number of floors and let $m = x_u y_u z_u$. Initially, the n facilities were placed at random in n out of the m cells where the bounds x_u,y_u,z_u were chosen such that m»n. One department at a time, say department D_k**,** was then lifted from its present position and placed in a cell C_k representing the optimal solution to "*weighted 1-median"*. The *weighted 1-median* problem was defined by the positions of the remaining n-1 facilities and with weights representing their communication with D_k. In case cell C_k already was taken up by another department, then the nearest cell among the m-n+1 free cells was chosen. This step was repeated for all n facilities taken in a random order. Afterwards followed an interchange procedure until a kind of "2-optimality" was achieved, that is, for the current layout of the n facilities, no further reduction of the objective function value CDIST would result from interchanging any pair of facilities.

Each run as described above resulted in a layout and a corresponding upper bound on the (unknown) minimum value $CDIST_{min}$ of CDIST. The series of such runs was continued until the relative difference between the best and the second best solution found was below a pre-specified threshold value or a certain time limit was exceeded.

For each instance of QAP examined, a set of the 10 best layouts found were presented to the contractor together with the corresponding values of CDIST. Based on these 10 values of CDIST, Spadille was furthermore requested to produce a good estimate of $CDIST_{min}$. Theoretical studies of QAP conducted much later would have confirmed Spadille's suspicion of a very flat optimum in general. Disregarding that suspicion, however, Spadille had no idea as to the real "landscape of optimal and near-optimal solutions". In hindsight, Krarup today is in no way proud of the statistical tools used to estimate $CDIST_{min}$. However, the contractor was satisfied and very pleased to see that no architect managed to come up with a layout beating Spadille's layout in terms of CDIST.

In addition to n and x_u, y_u, z_u, the input data for each run consisted of a symmetric nxn matrix $C=(c_{ij})$ specifying the communications between each pair of facilities. Finally, for a pair (D_r, D_s) of cells with coordinates (x_r, y_r, z_r) and (x_s, y_s, z_s) respectively, the distance d_{rs} between these was calculated as

$$d_{rs} = c_1 \times (|x_r - x_s| + |y_r - y_s|) + c_2 \times |z_r - z_s|$$

where c_1, c_2 are constants reflecting average transport + waiting time incurred by horizontal and vertical movements in the building. The great computational advantage of the *Manhattan distance measure* employed is that each execution of *weighted 1-median*, the core of the randomized heuristic, requires *0(n)* steps only.

3. Early Computational Results

For the Krarup 30a instance of QAP with n = 30, Spadille used $c_1 = 50$, $c_2 = 115$ in the distance function above. Krarup has no recollection of the exact values of x_u, y_u, and z_u. However, they were all "sufficiently large", perhaps as large as 10. The best facility layout found in 1972 had CDIST = 91,730. The most attractive structure had two levels, each arranged in a 4x4-cell grid, with two of the cells "free".

In 1976, Burkard and Stratmann (Burkard, 1978) managed to reduce CDIST to 90,420 by means of the *perturbation method* supplemented by other refinements, producing the following layout:

z=1				z=2			
28	10	23	*	12	22	15	*
13	7	21	29	11	5	4	6
9	8	20	24	17	16	2	3
14	27	25	19	18	1	26	30

In the resulting two stories high building (z=1, z=2), the numbers shown refer to the facilities and "*" designates a free cell.

Krarup looked at the cells marked by [ ] in the copy of the Burkard-Stratmann layout below:

z=1				z=2			
28	10	23	*	12	22	[15]	[*]
13	7	21	29	11	[5]	[4]	6
9	8	20	24	17	[16]	2	3
[14]	27	25	19	[18]	[1]	26	30

and then proposed a *cyclic rotation* of the marked cells,

z=1				z=2			
28	10	23	*	12	22	4	15
13	7	21	29	11	16	5	6
9	8	20	24	17	1	2	3
*	27	25	19	14	18	26	30

thereby achieving a layout with a value of CDIST = 89,100.

It was not realized at the time, but, a simple pair-wise interchange of facilities 1 and 17 would have given the following optimum layout, with CDIST = 88,900:

z=1				z=2			
28	10	23	*	12	22	4	15
13	7	21	29	11	16	5	6
9	8	20	24	1	17	2	3
*	27	25	19	14	18	26	30

This last example points out the potential effectiveness of local search as a post-processing technique. A simple "2-opt" (Stützle and Dorigo, 1999) local search would today be capable of finding the optimum solution. Though, this would have been difficult to implement with the computing power available in 1976.

It was at about this point that the floor plan devised in the above "cyclic rotation" became fixed and the corresponding facility flow and distance matrices entered into the QAPLIB library of QAP instances. All experimentation was thereafter made on this flow-distance model.

Not much happened regarding the Krarup 30a until the late 1980s. Then, an optimum objective function value of 88,900 was found by Jadranka Skorin-Kapov (Skorin-Kapov, 1990) using a "strict" tabu search. This result was published in 1990. As the tabu search did not guarantee optimality, the question remained. Is there a better solution?

4. Later Computational Results.

The Krarup 30a instance became a standard benchmark problem. During the ensuing decade, a number of heuristic approaches for solving the QAP were developed. Among these are Taillard's robust tabu search (Taillard, 1991), simulated annealing (Connolly, 1990), iterated local search (Stützle, 1999), greedy randomized adaptive search procedure (GRASP) (Li, 1994), scatter search (Cung,, 1997), and genetic algorithms (Fleurent, 1994), (Merz, 1997) and (Tate, 1995). The solution value 88,900 for the

Krarup 30a was re-found many different times. Taillard (Taillard, 1995) said that problem instances of this nature have "pseudo-optimal" solutions.

Then, in early 1999, Thomas Stützle did some experiments on the search space analysis with QAP instances. He examined the fitness-distance correlation (FDC) which captures the relationship between the quality of solutions (in that research of locally optimal solutions) and the distance to the closest globally optimal solution (Jones, 1995 and Boese, 1997). For the QAP, an appropriate distance measure between two solutions is the number of items that are placed on distinct locations.

When applying the FDC analysis to QAP instances, one is faced with the problem that for many QAP instances there exist a large number of globally optimal solutions. For example, if the distance matrix stems from a n_1 by n_2 grid and the distances are defined as the Manhattan distance between grid points, then the resulting QAP instances have multiple global optima (at least four if $n_1 \neq n_2$ and at least eight if $n_1 = n_2$ (Taillard, 1995). Even worse, these globally optimal solutions may be at the maximally possible distance from other globally optimal solutions.

For these reasons, Stützle had to find many "optimal solutions" and one of the QAP instances that he investigated was the Krarup 30a. For this instance he generated 257 "global optima", each with an objective function value of 88,900. The "global optima" were found using an Iterated Local Search algorithm (Stützle, 1999). After sorting, it was determined that 256 of these are distinct. For some of the other QAP instances listed in QAPLIB he found even many more "optimal" solutions. For example, for the bur26h and the bur26f instances he found more than 1000 "optimal" solutions. In all his experiments Stützle assumed that the solutions were optimal. In fact he always stopped at the best, known solution value. Thus, there was still no proof of the optimality of the 256 solutions that he had found for the Krarup 30a.

5. The Search for Optimality

At about the same time that Stützle did his work on Iterated Local Search, Hahn, et. al. (Hahn, 1999) were setting out to test their improved branch-and-bound algorithm on a relatively difficult QAP instance. They picked the Krarup 30a because it had not been solved exactly and because the best known objective function value was a greater distance from the best lower bound than for most instances of that size. This indicated its general difficulty and made it more likely that an optimum objective function value might not yet have been found.

Hahn had based his branch-and-bound technique on a long neglected lower bounding scheme that he had devised in his 1968 dissertation (Hahn, 1968). His research objective in 1968 was to optimize the assignment of binary code words in the transmission of digital signal values that represent a continuum of measurements. The idea was to assign exactly n code words to exactly n quantization levels in such a way that the mean-square error in the resulting output signal values due to transmission errors would be minimized. This turned out to be a QAP and thus he concerned himself with the more general problem.

The lower bounding scheme was called the Dual Procedure (DP). It had been devised in an attempt to solve the QAP using the concepts of the Hungarian method for the Linear Assignment Problem. The DP actually did solve QAP instances of size 7 and smaller, however it failed to solve the larger test problems. At the time, Hahn did not realize that theoretical limitations made this impossible, except in certain fortuitous instances. He recognized that the DP produced a good lower bound, but decided not to try branch-and-bound and instead pursued a cutting plane technique which worked sufficiently well to solve several QAP instances of size 8. His cutting plane technique was an extension of work by House, Nelson and Rado (House, 1965) on the resolvent sequence for solving 0-1 integer programs. The DP was used in the resolvent sequence approach as a means for fixing solution variables once the search space had been reduced.

Hahn's QAP research lay dormant until 1987 when a student, Thomas Grant decided to undertake the QAP as a Master's thesis (Grant, 1989) topic. Together, Hahn as advisor and Grant worked on the original cutting plane approach and used it to solve QAP test instances of size 12. It became clear, however, that the resolvent sequence was no match for the indomitable QAP. During this period Grant recognized the potential for the DP and recommended a branch-and-bound approach. He made some tests to verify that the DP was indeed better than any other lower bounding procedure existing at that time, with the exception of eigenvalue bounds which are not suitable in branch-and-bound. He also suggested several changes in the DP calculation to improve the bound value reached.

Again there was a lull in this line of research. Both advisor and student held full-time jobs and there were little resources for QAP research. But, in 1996, Hahn began to work on a branch-and-bound algorithm. He recognized immediately that certain advantages could be had from the DP that could not be had from other bound calculations. These are:

- Similar to a number of other lower bounding procedures, the DP generates a series of non-decreasing lower bounds on the QAP. More importantly, with each such lower bound, it generates a fully equivalent QAP with a reduced cost objective function. When branching takes place on such a reduced cost equivalent QAP, the ensuing bound calculations are performed more rapidly.
- Each lower bounding calculation at a node of the tree (i.e., for a given partial assignment) is an attempt to solve a reduced size QAP (i.e., a sub-problem of the original) from a dual perspective. In many cases, the DP solves sub-problems that are only a few sizes smaller than the QAP being tackled. Thus, feasible solutions to the original QAP are generated without having to fathom deep into the search space.
- The DP calculation is an iterative process that permits stopping early. It can be stopped as soon as the lower bound on the assumed partial assignment exceeds an upper bound on the original problem. It can also be stopped, in favor of making an additional partial assignment, when it becomes obvious that from its slow progress that it is unlikely to ever reach the upper bound. This is a very effective way to reduce branch-and-bound run time.

The branch-and-bound algorithm based on the DP worked quite well. Hahn, Grant and Hall (Hahn, 1998) tested it on problem instances from QAPLIB. During these tests, the single processor runtime for the algorithm was significantly shorter than that of the best known algorithms for instances of size 16 through 22. The runtime for the very difficult Nugent 22 instance was over 26 times faster than the only other competing algorithm at that time (Clausen, 1997). The Nugent problem instances are noted for their difficulty. Nugent, et. al. (Nugent, 1968) posed these problem instances of size 6,8,12,15,20 and30. Based upon the Nugent construction, researchers have added problem instances of size 18, 21, 22, 27 and 28 as intermediate benchmarks.

After submitting the branch-and-bound results for publication, Hahn and Grant (Hahn, 1998) set out to publish the test results from the latest version of the DP. It would later be referred to as the Hahn-Grant bound (HGB). With advice from Terry Johnson, Warren Adams and Monique Guignard-Spielberg, they were able to add the following important explanation: the DP and therefore the HGB procedure can be viewed in the context of linear programming. In fact, the DP approach is a dual-ascent procedure applied to a linear formulation of the QAP. It bears strong resemblance to work of Johnson [Joh92] and Adams and Johnson (Adams, 1994). The DP also bears resemblance to other existing lower bounding techniques, such as Assad and Xu (Assad, 1985), Burkard (Burkard, 1991), Carraresi and Malucelli (Carraresi, 1992), Gilmore (Gilmore, 1962), Lawler (Lawler, 1963), and Roucairol (Roucairol, 1979). The difference lies in the DP’s broad class of operations that permit more powerful and more computationally efficient strategies than do these others.

In 1998, there was little hope of solving problems much larger than the Nugent 22 in reasonable CPU time. More work had to be done to improve existing algorithms. During attempts to solve QAPs of size greater than 20, it was noted that some branches of the branch-and-bound tree were eliminated more rapidly than were others. This suggested that branching strategy could play a significant role. After some experimentation, it was determined that many of the techniques and strategies for branch-and-bound enumeration suggested in prior works (Burkard, 1991), (Clausen, 1997), (Bazaraa,, 1983), (Kaku, 1986), (Pardalos, 1989) and (Pierce, 1971) were not helpful. Thus, it became necessary to work out new and better branching schemes.

What seemed to show promise was the polytomic depth-first search strategy of Mautor and Roucairol (Mautor, 1994). The Mautor-Roucairol (M-R) strategy improved runtime for problem sizes up to N=15. Beyond that, the strategy failed to give an improvement. It was concluded that the measure used by M-R for a facility/location branching decision was too weak and that better measures were needed. The measure used by M-R was a constant linear cost that would be added to the objective function were the branch taken. But, this concept is just another way of estimating the lower bound on the branching decision. Thus, developing a branching strategy based upon lower bound calculations was the way to go.

In 1999, Hahn, Hightower and Johnson (Hahn, 1999) developed for purposes of tree elaboration, a series of bound estimates ranging from rather poor estimates that are computationally trivial to much, much better bound estimates that are computationally expensive. This was based upon the finding that the better estimates are required for speeding up the enumeration of larger QAP instances ($n>16$) while poorer estimates are good for smaller problem instances ($n\leq 16$). Two issues were important in this development. The first is the rule by which bound estimates influence branching. The second is the choice of bound utilized for the tree elaboration decision with respect to the level in the tree.

Regarding the first issue, it is necessary to select at each level of the tree a facility or location upon which to extend the elaboration. The bounds associated with all possible (as yet unassigned) facility/location pairs come into play here. If these bounds are arranged in a square matrix, where the rows represent facilities and the columns represent locations, the idea is to select the row or column where bounds are generally large. Experimental results indicated that the most effective rule is to choose the row or column whose average bound value is the greatest.

Regarding the second issue, they designed and implemented an increasingly improved set of bound calculations. The better of these bound calculations to be utilized near the root and the less accurate (poorer bounds) utilized further into the tree. The result is an effective and powerful technique for shortening the run times of problem instances in the range of size 16 to 25. Run times were decreased generally by five- or six-to-one and the number of nodes evaluated was decreased as much as 10-to-one. Later improvements in their strategy produced a better than 3-to-1 reduction in runtime so that overall runtime improvement was as high as 20-to-1 as compared to the earlier results published in [Hgh98]. This was what was needed to finally solve the Krarup 30a. exactly.

6. The exact solution

The branch-and-bound enumeration of the Krarup 30a was begun in December of 1998. The enumeration was started on a DEC Alpha 500 (300 MHz) single CPU machine. After only 1,346 minutes of processing, the data files were moved to a SUN Ultra 10 (360 MHz). Exactly, 12,208 nodes had been examined. The remaining enumeration was done entirely on the Sun Ultra 10, also with a single CPU. For various reasons, the enumeration run was stopped and started about 15 times. Before each stop a data dump file was produced which was used to restart the enumeration at the exact place of the last data dump. The resulting progress and runtimes were entered into an Excel spreadsheet.

On April 12, 1999 the enumeration was completed. The total run time was calculated to be 141,958.4 minutes where a slight correction was made to adjust for the slower speed of the DEC Alpha on the first 1,346 minutes of processing. This corresponded to a total runtime of 98.6 days. The number of nodes examined was 29,764,589, or about 7 nodes every 2 seconds. As expected, the optimum objective function value remained at 88,900. The number of major branches containing solutions was 4. The number of optimum solutions listed by program was 133 out of the 256 found by Stützle. This at first caused some consternation. However, the situation was easily explained.

Recall that the DP in certain cases actually solves the QAP upon which it is working. In fact, it does so every time it is working on a branch of the larger QAP that is under enumeration by the branch-and-bound algorithm. Recall, too, that in the Hungarian method for solving the linear assignment problem for an input of non-negative numbers, the problem is solved when in the cost matrix exactly n zero costs are found such that every row and every column contains a zero. At that point, several optimum solutions may be found at once. That is, several size n sets of zero cost elements may form distinct solutions. This is exactly what happened in the Krarup 30 instance. The algorithm discovered that a given branch contains a solution because the DP actually solved the sub-problem represented by that branch. The program looked for the first set of zero costs representing a solution and printed the solution to file. However, there were other solution(s) that existed and could also have been printed out, if necessary. This was not done in our implementation.

It should be noted that no one in the team looked into the origins of the Krarup 30a before deciding to try to solve it to optimality. The two authors of the present paper met for the first time at a conference in July of 1999. Without knowledge of the application, it was assumed that there were no symmetries that could be exploited in order to reduce the branch-and-bound search space. This was not the case. The distance matrix has four distinct symmetries. Had these been explicitly known to the enumeration program, the runtime would have been cut in at least half and much fewer nodes would have been needed.

7. What Next?

While great progress has been made on generating good solutions to large and difficult QAP instances, this has not been the case for finding exact solutions. In the late 1960s, it was an achievement to find the optimum solution to a difficult instance of size 8. In the 1970's and 80's, one could only expect to solve difficult instances for $n<16$. It was not until the mid-1990's that Clausen and Perregaard (Clausen, 1997) were able to enumerate the very difficult Nugent 20 instance. Clearly, much progress has been made since then. Still, the exact solving of the most difficult of the Nugent instances, the Nugent 30 is beyond the capability of reasonable computing resources. Current estimates for solving this instance using existing codes on a 360 MHz CPU range between 3 and 70 years.

Two approaches are promising. The first is that of the DP branch-and-bound technique which was used to solve the Krarup30a. The second is the newly developed branch-and-bound algorithm of Anstreicher and Brixius (Anstreicher, 2000) that uses a convex quadratic programming (QP) relaxation to obtain a bound at each node. In their method, The QP sub-problems are approximately solved using the Frank-Wolfe algorithm, which in this case requires the solution of a linear assignment problem on each iteration. Their branching strategy makes extensive use of dual information associated with the QP sub-problems.

Anstreicher and Brixius report state-of-the art computational results on large benchmark QAPs. For instance, on the Nugent 25 instance they achieved branch-and-bound enumeration evaluating 80,430,341 nodes in a wall time of 6.7 hours, with an average of 94 active 'workers'. The equivalent computation time on a single HP-C3000 workstation would be approximately 13 days. Whereas, the DP branch-and-bound enumeration took longer (52 days) on a single CPU Ultra 10 with a 360 MHz processor and required examination of fewer (38,726,326) nodes. Also, Anstreicher and Brixius are the first to have solved exactly the Nugent 27 instance. It remains to be seen which approach will perform better on even larger difficult instances. Have these approaches reached their potential? What new approaches are in the offing?

It has been argued that with the possibility of finding 'good' solutions quickly using heuristic methods such as simulated annealing, the search for exact solutions may not be worthwhile. This argument holds merit since we are dealing with models that are abstractions and not themselves reality. Therefore, it may not be necessary to have provably optimum solutions except in rather rare circumstances, when the model precisely matches the situation and large expenditures of time and money are involved. (This argument holds as well for metaheuristic techniques such as iterated local search, genetic algorithms and tabu search where the focus is on getting slight improvements over the simpler, faster heuristic methods.) A much better argument to support the search for exact solutions to difficult instances of QAP (and other intricate combinatorial optimization problems) is the permanent challenges of researchers and the hope of happy side effects in terms of devising new algorithmic tools or principles.

The QAP is a good problem to test new algorithmic ideas since it is so difficult. We can only ask that there will be renewed interest in the QAP and that more resources will be put at the disposal of those who are capable of making progress in this difficult area. This includes both enumerative strategies as well as heuristic strategies for finding very high quality solutions for larger instances. The race is on to solve the very difficult Nugent 30!

Acknowledgements

The authors thank Thomas Stützle for offering some valuable insights and for helping us to fill in the happenings between the origination of the Krarup 30 and its final solution.

References

Adams, W.P and Sherali, H.D, (1986) A tight linearization and an algorithm for zero-one quadratic programming problems, *Management Sci*ence, vol. 32, No. 10: 1274-1290.

Adams, W.P, and Johnson, T, (1994) Improved linear programming-based lower bounds for the quadratic assignment problem, *Quadratic Assignment and Related Problems*, Vol. 16 of *DIMACS Series on Discrete Mathematics and Theoretical Computer Science*: 43-76.

Anstreicher, K.M., and Brixius, N.W. (2000) Solving quadratic assignment problems using convex quadratic programming relaxations, University of Iowa Report, currently available on the Web at http://www.biz.uiowa.edu/faculty/anstreicher/qapqp2.ps

Assad, A. A., and Xu, W, (1985) On lower bounds for a class of quadratic 0,1 programs, *Operations Research Letters*, vol. 4: 175-180.

Bazaraa, M.S., and Kirca, O, (1983) A branch-and-bound-based heuristic for solving the quadratic assignment problem, *Naval Research Quarterly*, Vol. 30: 287-304.

Burkard, R. E., Karisch, S. E. and Rendl, F. (1991) QAPLIB - A quadratic assignment problem library, *European Journal of Operational Research*, vol. 55, no. 99: 115-119. This can be found on the web at http://serv1.imm.dtu.dk/~sk/qaplib

Burkard, R.E. and Strattman, K-H. (1978) Numerical investigations on quadratic assignment problems, *Naval Logistics Quarterly*, vol. 25, no. 1:129-148.

Boese, K.D. (1996) Models for iterative global optimization, *Ph.D. Thesis*, *University of California, Computer Science Department*, Los Angeles, USA.

Burkard, R.E. and Stratmann, K.H, (1978) Numerical investigations on quadratic assignment problems, *Naval Research Logistical Quarterly*, vol. 25, no. 1: 129-148.

Burkard, R. E. (1991) Locations with spatial interactions: The quadratic assignment problem, in *Discrete Location Theory*, P. B. Mirchandani and R. L. Francis, eds., John Wiley & Sons (1991): 387-437.

Carraresi, P, and Malucelli, F, (1992) A new lower bound for the quadratic assignment problem, *Operations Research*, vol. 40 Supple. no. 1: S22-S27.

Clausen, J. and Perregaard, M. (1997) Solving large quadratic assignment problems in parallel, *Computational Optimization and Applications*, 8:111-128.

Connolly, D. T. (1990) An improved annealing scheme for the QAP, *European Journal of Operational Research*, 46:93-100.

Cung, V-D, (1997) A scatter search based approach for the quadratic assignment problem, *Proceedings of ICEC'97*:165-170.

Fleurent, C. and Ferland, J. A. (1994) Genetic hybrids for the quadratic assignment problem, *Quadratic Assignment and Related Problems,* Vol. 16 of *DIMACS Series on Discrete Mathematics and Theoretical Computer Science*:173-187.

Frieze, A.M., and Yadegar, J. (1983) On the Quadratic Assignment Problem, *Discrete Applied Mathematics*, vol. 5, no. 1: 89-98.

Gilmore, P.C. (1962) Optimal and suboptimal algorithms for the quadratic assignment problem, *Journal of the Society of Industrial and Applied Mathematics*, vol. 10, no. 2: 305-313.

Grant, T.L. (1989) An evaluation and analysis of the resolvent sequence method for solving the quadratic assignment problem *Master's Thesis*, *University of Pennsylvania*,.

Hahn, P.M, (1968) Minimization of cost in assignment of codes to data transmission, *Ph.D. Dissertation*, University of Pennsylvania,.

Hahn, P.M., and Grant, T.L. (1998) Lower bounds for the quadratic assignment problem based upon a dual formulation, *Operations Research*, Vol. 46, No. 6: 912-922.

Hahn, P.M., Grant, T. L. and Hall, N. (1998) A branch-and-bound algorithm for the quadratic assignment problem based on the Hungarian method, *European Journal of Operational Research*, 108: 629-640.

Hahn, P.M., Hightower, W., Johnson, T., Guignard-Spielberg, M. and Roucairol, C. (1999) Tree elaboration strategies in branch and bound algorithms for solving the quadratic assignment problem, Submitted for publication to *Discrete Applied Mathematics*. Also, University of Pennsylvania Systems Engineering Dept. Working Report, currently available on the web at http://www.seas.upenn.edu/~hahn

House, R. W., Nelson, L. D., Rado, T. (1965) Computer studies of a certain class of linear integer programs, *Recent Advances in Optimization Techniques*, A. Lavi and T. P. Vogl, eds., John Wiley and Sons: 241-280.

Johnson, T.A, (1992) New linear programming-based solution procedures for the quadratic assignment problem, *Ph.D. Dissertation*, Clemson University, Clemson, SC.

Jones, T. and Forrest, S. (1995) Fitness distance correlation as a measure of problem difficulty for genetic algorithms, *Proceedings of the 6th International Conference on Genetic Algorithms*, Editor L. J. Eshelman, Morgan Kaufman, San Francisco: 184-19.

Kaku, B.K, and Thompson, G.L, (1986) An exact algorithm for the general assignment problem, *European Journal of Operational Research*, Vol. 23: 382-390.

Krarup, J. (1972) Quadratic Assignment, *Saertryk af data*, 3-72.

Krarup, J. and Pruzan, P.M., (1978) Computer-aided layout design, *Mathematical Programming Study 9*: 75-94.

Koopmans, T. C., and Beckmann, M. J. (1957) Assignment problems and the location of economic activities, *Econometrica*, vol. 25: 53-76.

Lawler, E. L. (1963) The quadratic assignment problem, Management Science, vol. 9: 586-599.

Li, Y., Pardalos, P. M. and Resende, M. G. (1994) A greedy randomized adaptive search procedure for the quadratic assignment problem, *DIMACS Series on Discrete Mathematics and Theoretical Computer Science*, Vol. 16: 237-261.

Mautor, T. and Roucairol, C. (1994) A new exact algorithm for the solution of quadratic assignment problems, *Discrete Applied Mathematics*, Vol. 55: 281-293.

Merz, P. and Freisleben, B. (1997) A genetic local search approach to the Quadratic Assignment Problem, *Proceedings of the Seventh International Conference on Genetic Algorithms (ICGA'97)*: 465-472.

Munkres, M. (1957) Algorithms for the assignment and transportation problems, *Journal of the Society of Industrial and Applied Mathematics*, vol. 5, no. 1: 32-38.

Nugent, C. E., Vollman, T. E and Ruml, J. (1968) An experimental comparison of techniques for the assignment of facilities to locations, *Operations Research*, vol.: 150-173.

Pardalos, P.M. and Crouse, J.V. (1989) A parallel algorithm for the quadratic assignment problem, *Proceedings of the 1989 Supercomputing Conference*, ACM Press: 351-360.

Pierce, J.F. and Crowston, W.B, (1971) Tree-search algorithms for quadratic assignment problems, *Naval Research Logistics Quarterly*, Vol. 18: 1-36.

Roucairol, C. (1979) A reduction method for the quadratic assignment problem, *Operations Research Verfahren-Methods of Operations Research*, vol. 32: 185-187.

Skorin-Kapov, J. (1990) Tabu search applied to the quadratic assignment problem. *ORSA Journal on Computing*, vol. 2, no. 1:33-45.

Stützle, T, (1997) MAX-MIN ant system for quadratic assignment problems, Research Report AIDA-97-04, Department of Computer Science, Darmstadt University of Technology, Germany.

Stützle, T. (1999) Iterated local search for the quadratic assignment problem, submitted for publication to *The European Journal of Operational Research*, also Research Report AIDA-99-03, Department of Computer Science, Darmstadt University of Technology, Germany,

Stützle, T. and Dorigo, M. (1999) ACO algorithms for the quadratic assignment problem, to appear in D. Corne, M. Dorigo and F. Glover, editors, *New Ideas in Optimization*, McGraw-Hill.

Taillard, E. D. (1991) Robust taboo search for the quadratic assignment problem, *Parallel Computing*, Vol. 17:443-455.

Taillard, E. D. (1995) Comparison of iterative searches for the quadratic assignment problem, *Location Science*, Vol. 3: 87-105.

Tate, D. M. and Smith, A. E. (1995) A genetic approach to the quadratic assignment problem, *Computers and Operations Research* 22 (1): 73-83.