Progress in Solving the Nugent Instances of the Quadratic Assignment Problem

Peter M. Hahn
The University of Pennsylvania
Peter_Hahn@CompuServe.Com
Department of Systems Engineering
293 Towne Building, 220 S. 33rd Street
Philadelphia, PA 19104-6315

1. The Quadratic Assignment Problem

The Quadratic Assignment Problem (QAP) is one in which N units have to be assigned to N sites in such a way that the cost of the assignment, depending on the distances between the sites and the flows between the units, is minimal. It can be formulated as follows:

Given two NxN matrices, $F = [f_{ij}]$ with f_{ij} the flow between units i and j, and $D = [d_{kl}]$ with d_{kl} the distance between sites k and l, find a permutation p of the set S = {1,2,...,N} which minimizes the global cost function, $\mathrm{Cost}(p) = \sum_{i=1,..,n} \sum_{j=1,..,n} f_{ij}\, d_{p(i)p(j)}$

The QAP is arguably the most difficult NP-hard combinatorial optimization problem. Solving general problems of size greater than 30 (i.e., with more than 900 (0-1) variables) is still computationally impractical. If among exact algorithms, branch-and-bound are the most successful ones, the lack of a sharp lower bound technique in these algorithms is one of the major difficulties.

The fact that the QAP is NP-hard is not sufficient to explain its difficulty, as we can now solve exactly very large instances of a great number of NP-hard problems. The homogeneity of the values of the solutions for most of the applications, due to the structure of the problem (scalar product of the two matrices) is a more convincing explanation. Indeed, first, we have a lot of solutions whose value is close to the optimum. So, even when the best solution is obtained, it is very hard to prove its optimality. Then, fixing one assignment has a low influence on the average value of the solutions. Even when going down in the branch-and-bound tree, the problem remains very hard. Moreover, it is difficult to prune important branches. Related to this aspect, the computation of the lower bound is one of the major difficulties. Indeed, the bound is either too loose (the number of nodes of the search tree becomes huge), or the time needed to compute the bound of a node is prohibitive.

2. The Nugent Instances of the QAP

In the late 1960's, Nugent, Vollmann and Ruml (NVR - 1968) posed a set of problem instances of size 5, 6, 7, 8, 12, 15, 20 and 30 noted for their difficulty. In these problems, the distance matrix stems from a n1 by n2 grid and the distances are defined as the Manhattan distance between grid points. The resulting QAP instances have multiple global optima (at least four if n1≠n2 and at least eight if n1=n2 (Taillard, 1995). Even worse, these globally optimal solutions are at the maximally possible distance from other globally optimal solutions.

In their paper NVR provided the exact solutions of the first four of these instances. However, they did not solve these with an efficient algorithm. They enumerated all the possible solutions to find the optimum. Enumeration time of the Nugent 8 was 3374 seconds on a GE 265 computer.

3. Early Computational Results (1970 - 1979)

A number of researchers set out to solve the Nugent instances using a variety of techniques. These included branch-and-bound, cutting planes and various heuristics. Burkard who solved the Nugent 8 in 0.426 seconds on a CDC-CYBER76 machine made the first notable progress in 1975. He used a branch-and-bound algorithm with Gilmore-Lawler (GL) lower bounds.

In 1978 Roucairol applied the reduction method of Little, Murty, Sweeney and Karel (LMSK 1963) to the QAP. She succeeded in finding a good sub-optimal solution to the Nugent 20 that was provably within 20% of optimum.

In 1978, Burkard and Stratmann optimally solved the Nugent 12 in 29.325 seconds on a CDC-CYBER76 machine. They used a branch-and-bound perturbation algorithm; utilizing the popular GL lower bound. The GL was to remain the preferred lower bounding technique until well into the 1990's.

Bazaraa and Sherali, in 1979, improved upon the computational results of Burkard and Stratmann. Using a formulation of concave quadratic programming and applying disjunctive cuts, they found the optimal solutions of the Nugent 12 in 21,8 seconds and of the Nugent 15 in 63.421 seconds on a CDC-CYBER74 machine. They did not prove optimality, A year later they applied Bender's partitioning scheme to a linearized objective function of the QAP and cut their runtimes by approximately half.

4. Later Computational Results (1980 - 1989)

In 1980 Burkard and Derigs reported that they were the first to solve the Nugent 15 using a branch-and-bound code. They did this in a runtime of 2947.32 seconds on a CDC-CYBER76. A few years later, Bazaraa and Kirka (1983), using a branch-and-bound-based heuristic found a solution of the most difficult of the Nugent instances, the 30, that is within 0.0007 of the best known solution today.

Then, in 1984, Picone and Wilhelm were able to produce all of the known best solutions of the eight original Nugent instances. They did this using a perturbation scheme to improve Hillier's earlier solutions. However, their paper on the subject contains a few apparent errors in the tabulated results.

Not much of consequence happened regarding the Nugent instances during the remainder of the decade, though research continued on the QAP in general.

5. Recent Computational Results (1990 - 1999)

In 1990 the Nugent instances were entered into the QAPLIB library of standard test instances and more researchers entered the contest for their solution. 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 values for the Nugent instances were re-found many different times. Taillard (Taillard, 1995) said that problem instances of this nature have "pseudo-optimal" solutions. This was also a decade which would demonstrate significant progress in the solution of the QAP.

Though the Nugent 15 had been solved years earlier by Burkard and Derigs, in 1994 it was solved again in only 121 seconds on a Cray 2 machine by Mautor and Roucairol using a new branch and bound algorithm. They also constructed a size 16 Nugent instance and proceeded to solve that on the Cray 2 in only 969 seconds.

A year later, Clausen and Perregaard using a parallel branch-and-bound algorithm on a 16 processor MEIKO computing surface with Intel i860 processors were the first to solve the

Nugent 20. This was to be the beginning of a new age in QAP successes. The wall runtime for solving the 20 was 57,963 seconds and the solution required the evaluation of 360,148,026 nodes. On a single processor the runtime would have been 811,440 seconds. For this, they relied upon the GL bound that had been developed decades earlier. Clausen and Perregaard also devised Nugent-like instances of size 14, 16, 17, 18, 21, 22, 24, and 25 to be used as intermediate benchmarks. These were constructed out of the larger instances by deleting certain rows and columns of the flow matrix and constructing an appropriate rectangular distance matrix. Of these, they solved the 14, 16, 17 and 18.

A very short time later, Hahn, Grant and Hall also solved the Nugent 20 using a unique branch-and-bound code based upon a dual procedure bounding technique that Hahn had devised in 1968 and which had never been published (outside of his dissertation). This was accomplished in 159,900 seconds on a single SPARC10 75MHz processor and required evaluation of only 724,289 nodes. At this point, results began to come more quickly. Brüngger, Marzetta, Clausen and Perregaard proceeded to solve the Nugent 21 and 22 as did Hahn, Grant and Hall. However, the latter team was able to solve these in single-processor-equivalent runtimes an order of magnitude faster and their solution required evaluation of far fewer nodes.

In 1997, Brüngger, Marzetta, Clausen and Perregaard continued in the lead by solving the Nugent 24 in 34 days on a 32 processor Parsytec CC machine with Motorola 604 processors. This was followed shortly by Hahn who solved the 24 in 56 days on a single 300 MHz processor DEC Alpha 500. A year later Brüngger and Marzetta proved the optimal solution to the Nugent 25. This took 23 days on a 128 processor Paragon XP/S22 -plus- 43 days on a GigaBooster of 7 Alpha 3000's -plus- 25 days on two more Alpha 3000s. A total of 276.6E9 nodes were evaluated in this exercise. While this solve the problem exactly, it is not considered a general solution since the approach was tailored for the 8-fold symmetries of the Nugent 25.

In 1998, there was little hope of solving problems much larger than the Nugent 24 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.

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 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 ($n\leq16$). This work permitted a general solution of the very difficult Nugent 25 instance. Hahn, Hightower and Johnson reported the general solution of the Nugent 25 early in 1998. It took 5,698,818 seconds on a single 360 MHz processor UltraSPARC 10 and required the evaluation of 108,738,131 nodes.

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, 1996). He used this technique to generate "optimal solutions" for several Nugent instances. For the Nugent 30 he generated 4 "global optima", each with an objective function value of 6124. These "global optima" were found using an Iterated Local Search algorithm (Stützle, 1999). 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 4 solutions that he had found for the Nugent 30.

6. Current Status

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 processor range between 3 and 70 years.

Two approaches are promising. The first is that of the DP branch-and-bound technique of Hahn, Grant and Hall. 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 of Hahn took longer (52 days) on a single CPU Ultra 10 with a 360 MHz processor but required examination of fewer (38,726,326) nodes.

Anstreicher and Brixius constructed Nugent 27 and 28 instances by removing the end rows and columns of the flow matrix of the Nugent 30. They used a 3 by 9 grid for the Nugent 27 distance matrix and a 4 by 7 grid for the Nugent 28 distance matrix. They are the first to have solved exactly these two instances. The Nugent 27 was solved in a wall time of approximately 24 hours on a Master-Worker (MW) distributed processing system developed as part of the MetaNEOS project at Argonne National Labs. During this 24 hours, the number of worker machines averaged 136 and peaked at 238. The equivalent computation time on a single HP-C3000 workstation would be approximately 126 days. It required the evaluation of 513,160,139 nodes. The Nugent 28 was solved in 4 days, 8 hours wall time. The number of active worker machines averaged approximately 200. The equivalent runtime on a single HP-C3000 workstation would be 435 days. This required the evaluation of 2,935,394,013 nodes.

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?

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!

References

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

Bazaraa, M. S. and Sherali, H. D. (1979) New Approaches for Solving the Quadratic Assignment Problem, ***Operations Research Verfahren-Methods of Operations Research***, vol 32: 29-46.

Bazaraa, M. S. and Sherali, H. D. (1980) Bender's Partitioning Scheme Applied to a New Formulation of the Quadratic Assignment Problem, Naval Research Logistics Quarterly, vol. 27: 29-42.

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.

Boese, K. D. (1996) Models for Iterative Global Optimization, ***Ph.D. Thesis***, University of California, Computer Science Department, Los Angeles, USA.

Burkard, R. E. (1975) Numerische Erfahungen Mit Summen - Und Bottleneck-Zuordnungsproblemen ***in Numerische Methoden Bei graphentheoretischen und kombinatorischen Problemen,*** Eds. L. Collatz and H. Werner, Birkhauser Verlag, Basel.

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. and Derigs, U. (1980) ***Assignment and Matching Problems: Solution Methods with FORTRAN Programs,*** New York, Springer-Verlag.

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.

Brüngger, A., Marzetta, A., Clausen, J and Perregaard, M. (1996) Joining Forces in Solving Large-Scale Quadratic Assignment Problems in Parallel, Institute for Theoretical Computer Science, Swiss Federal Institute of Technology, CH8092 Zurich, Switzerland.

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.

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.

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

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***, Ed. L. J. Eshelman, Morgan Kaufman, San Francisco: 184--192.

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.

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.

Little, J. D., Murty, K. G., Sweeney, D. W., Karel, C. (1963) An Algorithm for TSP***, Operations Research*** 11: 972-990.

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.

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.

Picone, C. J. and Wilhelm, W. E. (1984) A Perturbation Scheme to Improve Hillier's Solution to the Facilities Layout Problem, ***Management Science***, vol 30, no 10: 1238-1249.

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 quadratic assignment problems, ***Operations Research Verfahren-Methods of Operations Research***, vol. 32: 185-187.

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.