A Massively Parallel Solver for QAPs

Massively Parallel Solution of Challenging Large Scale Quadratic Assignment Problems

***Yuji Shinano**, Zuse Institute Berlin*

In Short

- The goal of the project is to solve previously unsolved challenging large scale Quadratic Assingment Problem (QAP) instances from QAPLIB.

For a positive integer n, we let $N = \{1, \ldots, n\}$ represent a set of locations and also a set of facilities. Given $n \times n$ symmetric matrices $\boldsymbol{A} = [a_{ik}]$, $\boldsymbol{B} = [b_{j\ell}]$ and an $n \times n$ matrix $\boldsymbol{C} = [c_{ij}]$, the quadratic assignment problem (QAP) is stated as

$$\min_{\pi} \sum_{i \in N}^{n} \sum_{k \in N}^{n} a_{ik} b_{\pi(i),\pi(k)} + \sum_{i \in N}^{n} c_{i,\pi(i)}, \qquad (0.1)$$

where a_{ik} denotes the flow between facilities i and k, $b_{j\ell}$ is the distance between locations j and ℓ, c_{ij} the fixed cost of assigning facility i to location j, and $(\pi(1), \ldots, \pi(n))$ a permutation of $1, \ldots, n$ such that $\pi(i) = j$ if facility i is assigned to location j. The QAP is one of the most famous classical combinatorial optimization problems. The QAPLIB [3] was first published in 1991 in order to provide a unified test bed for QAP that is widely accessible to the scientific community. Since its first publication, the QAPLIB has been updated continuously to stimulate further research of this important problem class. As a result, the QAP has continuously studied for the last three decades—theoretically and computationally.

In order to push the frontier of solvability of QAP, we need to combine the latest mathematical bounding techniques with the latest computational environments. Anstreicher et al. [1] solved the nug30 ($n = 30$) instance from QAPLIB for the fist time, by using new bounding techniques and new branch-and-bound schemes typically implemented in very powerful parallel computation environments (computational GRID) in 2000. At that time, the nug30 was solved over a period of seven days using 1000 machines. Recently, we solved the nug30 in an hour by using HPE SGI 8600 with 1,728 cores by using a newly develop parallel QAP solver based on UG [2]. In the technical report [2], we reported our intermediate research status to show possibility of advanced bounding techniques to push the frontier of solvability for QAP. To date, we have successfully solved tai30a and sko42 from QAPLIB for the first time. Our next target problem is QAPs with dimension at least 50.

Solving QAP instances of size larger than 35 in practice still remains challenging. Various heuristic methods for the QAP such as tabu search [9,10], genetic algorithms and simulated annealing have been developed. Those methods frequently attain near-optimal solutions, which often also happen to be the exact optimal solutions. The exactness is, however, not guaranteed in general.

Most existing methods for finding the exact solutions of QAPs are designed using a branch-and-bound (B&B) algorithmic framework [1,4]. As its name indicates, branching and (lower) bounding are the main procedures of the method. The bounding based on doubly nonnegative (DNN) relaxations has recently attracted a great deal of attention as it provides tight global bounds. In this research proposal, four different methods will be employed to accelerate the solving of challenging large-scale DNN relaxation problems—improving the overall solution performance for the QAP.

First two of the methods are SDPNAL+ [12] (a majorized semismooth Newton-CG augmented Lagrangian method for semidefinite programming with nonnegative constraints) and BBCPOP [5] (a bisection and projection method for Binary, Box and Complementary constrained Polynomial Optimization Problems). Some numerical results on these two methods applied to QAP instances with dimensions $n = 15$ to 50 from QAPLIB were reported in [5] , where BBCPOP attained tighter lower bounds for many of instances with dimensions 30 to 50 in less execution time. In addition, new and improved lower bounds were computed by Mittelmann using BBCPOP for the unknown minimum values of larger scale QAP instances, including tai100a and tai100b, see QAPLIB [3]. The third method is an alternating direction method of multipliers (ADMM) proposed by Oliveira et al. [7] in combination with facial reduction for robustness. The fourth method is the Newton-bracketing (NB) method [6], which was developed to further improve the lower bounds obtained from BBCPOP by incorporating the Newton method into BBCPOP.

Numerical results on the NB method and BBCPOP are given in Table 2 of [6], and ones on ADMM in Table 7 of [7]. Lower bounds for the QAP instances sko42, sko49, sko56, sko64, sko72, sko81, tai60a, and tai80a are routinely computed using these methods. While the lower bounds obtained by the NB method for the first three instances are not smaller than those by ADMM,the NB method clearly attained the tightest lower bound among the three methods for the last five larger-scale instances with dimension $n \geq 60$.

The main motivation of our project is to challenge larger scale QAP instances from QAPLIB [3] that have not been solved yet. We implement the NB method [6] combined with the B&B method in the specialized Ubiquity Generator (UG) framework [11] to find the exact solutions of large scale QAP instances. See Sections 2, 3, 4, 5 in [2] for precise and detailed description about the proposed algorithms.

UG is a generic framework to parallelize branch-and-bound based solvers and has achieved large-scale MPI parallelism with 80,000 cores [8]. We have developed parallel QAP solver based on UG. We solved challenging large scale instances, nug30, tai30a, tai35b, tai40b and sko42 on the ISM (Institute of Statistical Mathematics) supercomputer HPE SGI 8600, which is a liquid cooled, tray-based, high-density clustered computer system. The ISM supercomputer has 384 computing nodes and each node has two Intel Xeon Gold 6154 3.0GHz CPUs (36 cores) and 384GB of memory. All of the instances were solved as a single job. Table 1 shows the computational results. Note that tai30a and sko42 were solved to the optimality for the first time, which had remained unsolved for more than 20 years.

Table 1: *Numerical results on challenging large scale QAP instances.*

QAP instance	Opt.val	No. of nodes generated	Total execution time(sec) in para.	No. of CPU cores used
nug30	6,124	26,181	3.14e3	1,728
tai30a	1,818,146	34,000,579	5.81e5	1,728
tai35b	283,315,445	2,620,547	2.49e5	1,728
tai40b	637,250,948	278,465	1.05e5	1,728
sko42	15,812	6,019,419	5.12e5	5,184

The goal of this project is to (further) develop the parallel QAP solver based UG to handle over 100,000 cores to solve notoriously hard QAP instances to optimality. Several branching schemes for the algorithms and a problem specific ramp-up, which is a phase until all cores busy, mechanism is proposed in [2]. However, these features are currently not implemented nor adequately tested. After implementing and testing all of the proposed features, the results are to be published in peer-reviewed international journals. We believe that this project would serve as a flagship for the efficient solution of $\mathcal{NP}$-hard combinatorial optimization problems by using supercomputers.

WWW

https://www.zib.de/members/shinano, https://ug.zib.de/

More Information

[1] K. Anstreicher, N. Brixius, G. J-P., Linderoth, and J., *Solving large quadratic assignment problems on computational grids*, Math. Program., 91 (2002), pp. 563–588.

[2] K. Fujii, N. Ito, S. Kim, M. Kojima, Y. Shinano, and K.-C. Toh, *Solving challenging large scale qaps*, Tech. Rep. 21-02, ZIB, Takustr. 7, 14195 Berlin, 2021.

[3] P. Hahn and M. Anjos, *QAPLIB – A quadratic assignment problem library*. http://www.seas.upenn.edu/qaplib.

[4] P. Hahn, A. Roth, and M. Saltzman, M. abd Guignard, *Memory-aware parallelized RLT3 for solving quadratic assignment problems*, Tech. Rep. Optimization Online, 2013.

[5] N. Ito, S. Kim, M. Kojima, A. Takeda, and K. Toh, *BBCPOP: A sparse doubly nonnegative relaxation of polynomial optimization problems with binary, box and complementarity constraints*, ACM Trans. Math. Soft., 45 (2019).

[6] S. Kim, M. Kojima, and K. Toh, *A Newton-bracketing method for a simple conic optimization problem*, To appear in Optim. Methods and Softw.

[7] D. Oliveira, H. Wolkowicz, and Y. Xu, *ADMM for the SDP relaxation of the QAP*, Math. Program. Comput., 10 (2018), pp. 631–658.

[8] Y. Shinano, T. Achterberg, T. Berthold, S. Heinz, and T. Koch, *Solving open mip instances with parascip on supercomputers using up to 80,000 cores*, 2016 IEEE International Parallel and Distributed Processing Symposium (IPDPS), IEEE, 2016, pp. 770–779.

[9] J. Skorin-Kapov, *Tabu search applied to the quadratic assignment problem*, ORSA J. Comput., 2 (1990), pp. 33–45.

[10] E. Tailard, *Robust taboo search for the quadratic assignment problem*, Parallel Comput., 17 (1991), pp. 443–455.

[11] *UG: Ubiquity Generator framework*. http://ug.zib.de/.

[12] L. Q. Yang, D. F. Sun, and K. C. Toh, *SDPNAL+: A majorized semismooth Newton-CG augmented Lagrangian method for semidefinite programming with nonnegative constraints*, Math. Prog. Comp., 7 (2015), pp. 331–366.

Funding

BMBF Research Campus MODAL