

# On solving a hard quadratic 3-dimensional assignment problem

Hans D. Mittelmann · Domenico Salvagnin

Received: 29 September 2013 / Accepted: 19 December 2014 / Published online: 22 January 2015  
© Springer-Verlag Berlin Heidelberg and The Mathematical Programming Society 2015

**Abstract** We address the solution of a very challenging (and previously unsolved) instance of the quadratic 3-dimensional assignment problem, arising in digital wireless communications. The paper describes the techniques developed to solve this instance to optimality, from the choice of an appropriate mixed-integer programming formulation, to cutting planes and symmetry handling. Using these techniques we were able to solve the target instance with moderate computational effort (2.5 million nodes and 1 week of computations on a standard PC).

**Keywords** Mixed-integer programming · Quadratic assignment · Symmetry

**Mathematics Subject Classification** 90C11 Mixed integer programming · 90C90 Applications of mathematical programming · 90C57 Polyhedral combinatorics, branch-and-bound, branch-and-cut

## 1 Introduction

The axial quadratic 3-dimensional assignment problem (Q3AP) is a generalization of the standard quadratic assignment problem (QAP). In its general form, the problem is defined by a 6-dimensional matrix of costs  $c_{ijkpqr}$  of  $n^6$  coefficients, and can be formulated nonlinearly as

---

H. D. Mittelmann  
School of Mathematical and Statistical Sciences, Arizona State University,  
Tempe, USA  
e-mail: mittelmann@asu.edu

D. Salvagnin (✉)  
DEI, University of Padova, Padua, Italy  
e-mail: salvagni@dei.unipd.it

$$\min \sum_{i=1}^n \sum_{j=1}^n \sum_{k=1}^n \sum_{p=1}^n \sum_{q=1}^n \sum_{r=1}^n c_{ijkpqr} x_{ijk} x_{pqr} \quad (1)$$

$$\sum_{i=1}^n \sum_{j=1}^n x_{ijk} = 1 \quad \forall k = 1, \dots, n \quad (2)$$

$$\sum_{i=1}^n \sum_{k=1}^n x_{ijk} = 1 \quad \forall j = 1, \dots, n \quad (3)$$

$$\sum_{j=1}^n \sum_{k=1}^n x_{ijk} = 1 \quad \forall i = 1, \dots, n \quad (4)$$

$$x_{ijk} \in \{0, 1\} \quad (5)$$

Equivalently, the objective is to minimize the quadratic objective function (1) over the 3-dimensional assignment polytope defined by (2)–(5). Clearly, the problem is NP-hard, being a generalization of QAP and of the 3-dimensional linear assignment problem (3AP), both of which are NP-hard. Q3AP was introduced almost 50 years ago in a technical memorandum later published in [19], without further development in the literature, and it was recently rediscovered as a tool to model symbol remapping in digital wireless communications [7]. In this context, binary data is mapped to symbols for transmission and, in case of errors, some data packets are retransmitted, following an automatic repeat request (ARQ) protocol. Recent works [23, 24] have shown that during retransmission a different encoding should be used, in order to improve the probability of error detection. Intuitively, symbol mapping should be as diversified as possible between retransmissions. When only one retransmission is needed, the problem is equivalent to the solution of a standard QAP, see [26] for a derivation. However, if two retransmission are needed, then the problem is a Q3AP (provided that one wants to optimize the simultaneous assignment of binary strings to symbols in all transmissions). In this case, the cost coefficients  $c_{ijkpqr}$  represent the costs of assigning to the strings  $i$  and  $p$  the symbols  $j$  and  $q$  in the first retransmission and the symbols  $k$  and  $r$  in the second retransmission. These coefficients can be computed offline based on a probabilistic error model, see [23, 24] for details, and can lead to massive improvements in the corresponding communication quality. Note that while a Q3AP of size  $n$  has  $n! \times n!$  feasible solutions, the optimization has no real time requirements, because the problem needs to be solved in the planning and design phase, and not by the parties involved in the communication. Still, the Q3AP turns out to be an extremely difficult problem to solve: indeed, the exact methods in [7] can solve only artificially generated instances of size up to 13, even allowing for massive computing power.

The purpose of this paper is to describe our efforts in the solution of the previously unsolved Q3AP real world instance of size 16. In order to successfully solve this instance, we developed a few interesting techniques that may turn out to be instrumental for the solution of other Q3AP (or QAP) instances as well. The proposed techniques yield a solution procedure which is exact in nature: we decompose the original instance into subproblems based on symmetry arguments (which are symbolic, and thus exact

in nature), and then we formulate each subproblem as a mixed-integer programming (MIP) program, to be solved with a complete method such as branch-and-cut. However, since each subproblem is solved by a state-of-the-art, but finite precision, MIP solver, the optimality proof obtained by the implemented procedure depends on all code being correct and all floating point computations being done to sufficient accuracy. In the following, by optimality proof we will refer to this computational connotation, highlighting when needed all the precautions that we took in order to increase the numerical precision and fault tolerance of our procedure.

The outline of the paper is as follows: Sect. 2 describes the real world instance that we wanted to solve, while Sects. 6.2–6.4 present the techniques that we developed for the purpose, namely a lightweight mixed-integer programming formulation (Sect. 6.2), symmetry handling (Sect. 6.3) and cutting plane generation (Sect. 6.4). Section 6 reports the computational experiments. Finally, conclusions are drawn in Sect. 7.

## 2 Instance

There are several digital modulation techniques, such as ASK, FSK, PSK, QAM. For phase-shift keying (PSK) the carrier frequency of the modulated signal is constant and the modulation alphabet is represented in a constellation diagram. For an  $n$  symbol PSK or  $n$ -PSK, this shows  $n$  points distributed on the unit circle equidistantly with respect to angle. This is a simple pattern and one that leads to a high degree of symmetry which allows to apply and further refine symmetry exploiting methods that have lately been developed for general mixed-integer programs. The data for this real world instance is available in [6].

Regarding the exact solution of the Q3APs associated with ARQ modulation techniques, so far only the 8-PSK case has been solved. For  $n = 16$ , only the successive optimization of the symbol mappings, which leads to several standard QAPs, has been considered in [24], while in this paper we tackle the simultaneous optimization of all mappings. The cost coefficients for the 16-PSK case are, in general, floating point numbers. As is customary, the coefficients were scaled by  $10^{16}$  and rounded to integers: this introduces an approximation error that is orders of magnitude smaller than the default tolerances used by any commercial MIP solver, so it is absolutely safe in our context. From now on, we will refer to this all-integer version as our instance to solve. Table 1 lists some of its characteristics, such as objective density and dynamism<sup>1</sup>, before and after scaling (see Sect. 6.2 for the details about our scaling procedure). Clearly, scaling improves sparsity and objective dynamism significantly. While we show in the present paper that it is possible to solve the Q3AP derived from 16-PSK, bigger instances such as 32-QAM, 64-QAM, etc can only be treated with a reasonable effort by heuristic methods to obtain feasible solutions, combined with lower-bounding techniques to show which gap in objective value is present. For 2D index assignment problems up to dimension 512 this was done in [26], based on methods developed in [15, 18]. The approach can be generalized to Q3AP problems.

---

<sup>1</sup> The dynamism of the objective function is defined as the ratio between the maximum and minimum absolute values of the non-zero coefficients.

**Table 1** Instance characteristics

$n$	16
$n^3$	4,096
Non-zero objective coefficients	12,755,712
Objective sparsity	76.03 %
Objective dynamism	$3.627 \times 10^{12}$
Symmetry group order	49,152
Non-zero scaled objective coefficients	7,524,096
Scaled objective sparsity	44.85 %
Scaled objective dynamism	$3.627 \times 10^6$

### 3 MIP Model

In order to solve such a large Q3AP instance with MIP technology, a lightweight, yet sufficiently strong, formulation is needed. A trivial way to linearize the quadratic terms in the objective function is to introduce additional binary variables  $\psi_{ijkpqr} = x_{ijk}x_{pqr}$ , and the corresponding linking constraints. However, in our case this amounts to introducing  $\Theta(n^6)$  variables and constraints, which is hopeless for  $n = 16$ . To overcome this issue, we preferred to extend the lightweight QAP Kaufman and Broeckx (KB) model [9] to the Q3AP case. The model requires only  $n^3$  artificial continuous variables, defined as

$$w_{ijk} = \left( \sum_{p=1}^n \sum_{q=1}^n \sum_{r=1}^n c_{ijkpqr} x_{pqr} \right) x_{ijk}$$

which can be easily linearized with big-M coefficients. The corresponding MIP model reads

$$\min \sum_{i=1}^n \sum_{j=1}^n \sum_{k=1}^n w_{ijk} \tag{6}$$

$$\sum_{i=1}^n \sum_{j=1}^n x_{ijk} = 1 \quad \forall k \in \{1, \dots, n\} \tag{7}$$

$$\sum_{i=1}^n \sum_{k=1}^n x_{ijk} = 1 \quad \forall j \in \{1, \dots, n\} \tag{8}$$

$$\sum_{j=1}^n \sum_{k=1}^n x_{ijk} = 1 \quad \forall i \in \{1, \dots, n\} \tag{9}$$

$$w_{ijk} \geq \sum_{p=1}^n \sum_{q=1}^n \sum_{r=1}^n c_{ijkpqr} x_{pqr} - M(1 - x_{ijk}) \quad \forall (i, j, k) \in \{1, \dots, n\}^3 \tag{10}$$

$$x_{ijk} \in \{0, 1\} \tag{11}$$

$$w_{ijk} \geq 0 \tag{12}$$

Note that if all cost coefficients are integer, as in our case, artificial variables  $w_{ijk}$  could be declared integer as well, yielding a pure integer formulation. This model is known (in its QAP version) to be of little use in practice because of the big-M constraints (10). In particular, it can be proved [27] that the root-node bound is always zero. However, a much stronger formulation can be obtained by adding to (6)–(12) the (polynomial) family of cutting planes

$$w_{ijk} \geq L_{ijk}x_{ijk} \tag{13}$$

where each  $L_{ijk}$  is defined as the optimal value of the 3AP

$$\min \sum_{p=1}^n \sum_{q=1}^n \sum_{r=1}^n c_{ijkpqr}x_{pqr} \tag{14}$$

$$\sum_{p=1}^n \sum_{q=1}^n x_{pqr} = 1 \quad \forall r \in \{1, \dots, n\} \tag{15}$$

$$\sum_{p=1}^n \sum_{r=1}^n x_{pqr} = 1 \quad \forall q \in \{1, \dots, n\} \tag{16}$$

$$\sum_{q=1}^n \sum_{r=1}^n x_{pqr} = 1 \quad \forall p \in \{1, \dots, n\} \tag{17}$$

$$x_{ijk} = 1 \tag{18}$$

$$x_{pqr} \in \{0, 1\} \tag{19}$$

Note that, contrary to the QAP case [3,27], in order to compute the coefficients of this family of cutting planes, we need to solve a 3AP for each cut, which is by itself an NP-hard problem. Still, for the instance at hand, those 3AP instances were straightforward to solve with state-of-the-art MIP technology. In addition, the symmetry of the instance can be exploited to solve much fewer than  $n^3$  3APs (see Sect. 6.3). The KB model, together with the family of cutting planes (13), was used recently in [3] to solve highly symmetric QAP instances, for the same reasons. Finally, we note that the big-M coefficients in constraints (10) can also be computed by solving 3APs akin to the model just described, with the only differences that maximization is used instead of minimization, and the fixing  $x_{ijk} = 1$  is replaced with  $x_{ijk} = 0$ . While these 3APs turn out to be no harder to solve in practice than those needed to compute coefficients  $L_{ijk}$ , and the same symmetry trick applies, their computation is not even needed if constraints (10) are modeled as indicator constraints, a feature implemented in most commercial MIP solvers.

Formulating the problem as a MIP and trying to solve it with a black box MIP solver is however far from sufficient to solve the instance to optimality. For reference, we

**Table 2** Orbit characteristics

Orbit	Size	$L$	bigM
0	256	0	6.4325e+06
1	1,024	2	6.8216e+06
2	512	1,199	6.8349e+06
3	1,024	30	9.2371e+06
4	1,024	3,544	9.6247e+06
5	256	46,465	1.2042e+07

report that a default CPLEX run on model (6)–(12), with the global cutting planes (13) added at the very beginning and the optimal solution provided as incumbent, yields a depressing relative gap of 94.25 % left to close after 3 days of computations, more than 35 millions of enumerated nodes and 120 GB of memory occupied.

#### 4 Symmetry handling

Given an arbitrary MIP  $P$ , the formulation (or symmetry) group  $G_P$  of  $P$  is defined as the group containing the set of permutations  $\pi \in S_n$  that leave the formulation of  $P$  unchanged, except for a possible reordering of the constraints. The formulation group of  $P$  can be computed from its formulation using standard graph automorphism packages, such as, e.g., Nauty [14], see [10] for details. We applied such an algorithm to the model (6)–(12), and found a non trivial formulation group of order 49,152. While not an extreme case of symmetry by any measure, this group is still large enough to slow down the search significantly, in particular given that each branch-and-cut node is quite expensive in our case. According to the symmetry group, the set of binary variables  $x_{ijk}$  can be partitioned into 6 orbits (the same applies to the artificial variables  $w$ ), whose main figures are reported in Table 2. As anticipated in the previous section, we can exploit this knowledge to significantly reduce the effort needed to compute coefficients  $L_{ijk}$  and the big-Ms in the model, given that they will clearly be the same for all the variables in the same orbit. This reduces the number of 3APs to solve from 8,192 down to 12. Note that there is no chicken-and-egg problem here, because the symmetry group of the model can be computed *before* knowing the values of the coefficients [cuts (13) are not yet in the model, and all big-Ms can be set to the same value].

Symmetry has long been recognized as a curse for the traditional enumeration approaches used in many optimization communities—we refer to [5,13] for recent surveys on the subject. Various techniques for dealing with symmetric problems have been studied, with *isomorphism pruning* [12] and *orbital branching* [17] being the most popular and effective strategies deployed within MIP solvers. In a preliminary computational study, we implemented both symmetry breaking techniques: however, their effect was still not sufficient to solve our instance to optimality.

A different approach to symmetry management has been proposed recently and is called *orbital shrinking* [2]: the main idea is to aggregate variables along orbits, in order to obtain a smaller (and symmetry free) reformulation. While this approach can

yield an exact reformulation in some cases, in general it provides only a relaxation of the original problem: however, orbital shrinking has been turned into an exact method in [21]. The algorithm in [21] is a decomposition scheme, where a *master* problem is used to enumerate the feasible solutions of the orbital shrinking reformulation of the model, and a *slave* problem is solved for each such solution to check if it can indeed be turned into a solution of the original problem (see [21] for details). Orbital shrinking has been shown to be a successful approach in solving symmetric optimization problems with the *right* amount of symmetry [21, 22]: intuitively, for the method to be effective, the shrunk reformulation must be sufficiently easy to solve, yet retain enough of the structure of the original model. As such, it is a natural candidate for our case, where the amount of symmetry is not extreme.

The exact method in [21] can, in principle, be applied to pure-integer as well as mixed-integer problems: however, in the mixed-integer case, one should enumerate, in the master, all possible values also for the continuous variables, which makes the method impractical. Even in our case, where as noted we could declare variables  $w_{ijk}$  as integer, the number of solutions to check would be enormous. Thus, in order to apply orbital shrinking, we had to modify the method in the following way. Given the orbital shrinking reformulation of the model (which is obtained by aggregating variables and constraints along orbits, and, by definition, has one variable for each variable orbit and one constraint for each constraint orbit), we removed the aggregated constraints corresponding to (10), the aggregated continuous variables, and the objective function. The final model (which by construction is still a relaxation) reads:

$$y_0 + y_1 + y_2 + y_3 + y_4 + y_5 = 16 \tag{20}$$

$$y_1 + 2y_3 + y_4 = 16 \tag{21}$$

$$y_2 + y_4 + 2y_5 = 8 \tag{22}$$

$$2y_0 + y_1 + y_2 = 8 \tag{23}$$

$$y_i \in \{0, \dots, |O_i| - 1\} \quad \forall i \in \{0, \dots, 5\} \tag{24}$$

where we denoted with  $y_i$  the sum of all the binary variables  $x_{ijk}$  in orbit  $O_i$ , i.e.

$$y_i = \sum_{(i,j,k) \in O_i} x_{ijk}$$

Model (20)–(24) is of no use as far as the bound is concerned; however, it can be used to enumerate all possible 6-tuples of  $y$  variables. Each 6-tuple  $\bar{y}^t$  can then be used to construct a subproblem in which the sums of the variables along the orbits are fixed to  $\bar{y}^t$ . It is easy to see that this is a partition of the feasible space of the original problem. Luckily enough, in our case there are only 85 different 6-tuples (reported in Table 3), which implies that we need to solve *only* 85 subproblems. Similar enumeration ideas were used also in [11, 16].

It is important to note that each subproblem, by construction, has the same symmetry group as the original problem, a property shared with the exact algorithm in [21]. Usually, it is possible to get around this issue by exploiting ad-hoc symmetry breaking procedures for the subproblem at hand: however, in our case, we just resorted to

traditional general purpose techniques such as orbital branching and isomorphism pruning. To sum up, the proposed strategy exploits the symmetry group of the original model twice:

1. computation of the orbital shrinking inspired model (20)–(24)
2. solution of each subproblem with isomorphism pruning

Finally, we note that the constraints used to fix the values of the orbits in each subproblem can be taken into account when computing the coefficients  $L_{ijk}$ , resulting in an additional strengthening of the formulation.

## 5 Cutting planes

Cutting planes (13) turn out to be crucial for the strength of the KB formulation at the root node. Still, the resulting formulation is not strong enough for a pure branch-and-bound to succeed. If we dig deeper into the structure of this family of cutting planes, we can clearly see the weak spot: each  $L_{ijk}$  is obtained by solving, independently, a linear 3AP problem, thus the corresponding assignments are not *synchronized* among the  $n^3$  subproblems. Indeed, if all assignments were the same, then the KB model would give the value of the optimal solution: unfortunately, this is, by far, not the case. A possible way to (partly) fix this issue is to keep on separating local versions of cuts (13) throughout the search tree, using the current domain of the variables for the 3APs. This is beneficial because, at a given node of the tree, all the assignments must at least agree on the variables that are fixed at that node. Of course, this also means that all cutting planes generated this way are valid only locally (i.e., only within the subtree rooted at the node in which they are separated), so some bookkeeping is needed in order to guarantee correctness. The same strategy was also exploited in [3], where, however, each cut could be separated much faster by an ad-hoc polynomial procedure, much in the spirit of the classical Gilmore-Lawler bound computations [1], while in our case each 3AP is solved by a black box MIP solver.

While computational experiments show that separating local cuts of the family (13) is indeed very beneficial for the resulting enumerative search, still this is not enough to solve our instance. This is easily explained by the fact that we need to dive deep into the tree for the local domains to significantly constrain the separation problems. Intuitively, fixing a few variables does not yield sufficiently strong cuts, and the dual bound is still too weak to prune nodes early on.

Thus, we devised a new family of cutting planes, which is a generalization of (13). While a given cutting plane of the form (13) provides a lower bound on the value of a variable  $w_{ijk}$ , given the current domains of variables  $x$  and assuming that  $x_{ijk} = 1$ , the new family gives lower bounds on the sums of the form  $w_{ijk} + w_{pqr}$ , assuming that  $x_{ijk} = 1$  and  $x_{pqr} = 1$ . The resulting cuts are all of the form:

$$w_{ijk} + w_{pqr} \geq T_{ijkpqr}(x_{ijk} + x_{pqr} - 1) \quad (25)$$

and the coefficient  $T_{ijkpqr}$  can be obtained as the optimal value of the following MIP:



$$\min w_{ijk} + w_{pqr} \tag{26}$$

$$\sum_{s=1}^n \sum_{t=1}^n x_{pqr} = 1 \quad \forall u \in \{1, \dots, n\} \tag{27}$$

$$\sum_{s=1}^n \sum_{u=1}^n x_{pqr} = 1 \quad \forall t \in \{1, \dots, n\} \tag{28}$$

$$\sum_{t=1}^n \sum_{u=1}^n x_{pqr} = 1 \quad \forall s \in \{1, \dots, n\} \tag{29}$$

$$w_{ijk} \geq \sum_{s=1}^n \sum_{t=1}^n \sum_{u=1}^n c_{ijkstu} x_{stu} \tag{30}$$

$$w_{pqr} \geq \sum_{s=1}^n \sum_{t=1}^n \sum_{u=1}^n c_{pqrstu} x_{stu} \tag{31}$$

$$x_{ijk} = 1 \tag{32}$$

$$x_{pqr} = 1 \tag{33}$$

$$x_{stu} \in \{0, 1\} \tag{34}$$

Note that by construction  $T_{ijkpqr} \geq L_{ijk} + L_{pqr}$ . In our computational experience, each cut of the family (25) is not significantly harder to separate than a cut (13), yet care must be taken because the number of possible cuts at each node is  $n^6$ , which, although polynomial, could make an exact separation too slow in practice. Finally, note that, as in the previous section, the constraints used to fix the values of the orbits in each subproblem can be taken into account when computing the coefficients of cuts (13) and (25), both globally and locally.

## 6 Computations

In this section, we shall describe the tuning of the code which was necessary to perfect the techniques presented above.

### 6.1 Primal heuristics

After the first runs on the complete model (6)–(12), we realized that even a state-of-the-art solver such as CPLEX has lots of trouble in finding a good quality (hopefully optimal) feasible solution early on in the search. This is of course affecting negatively performance in general, and the more so in our case where the LP relaxation is quite weak. In addition, this is a huge issue for the decomposition scheme: not knowing the globally optimal solution may turn some easy subproblems into very difficult ones, and carrying over the best solution from one subproblem to the next is not an option (we would not be able to process subproblems in parallel, and in any case the sequential order would be arbitrary).

In order to overcome this issue, we implemented an ad-hoc primal heuristic for the Q3AP, along the lines of the ILS method described in [7, 25]. Surprisingly, even a very simple implementation was able to consistently find the (later proven) optimal solution of value 207, 462, 238, 240 in a matter of minutes. As such, we can, for practical purposes, assume that we know the optimal solution in advance, and concentrate on the task of proving its optimality. For reference, we provide the optimal solution below:

```

0 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15
0 1 2 3 4 6 5 7 8 10 9 11 12 13 14 15
6 14 2 10 8 12 0 4 11 15 3 7 5 13 1 9

```

## 6.2 MIP Model

A first issue is given by model (6)–(12) itself: because the objective 6-dimensional matrix is quite dense, the corresponding constraints (10) are dense as well. In addition, the dynamism of these coefficients is so high (values range from 1 to  $3.627 \times 10^{12}$ ) that the LP relaxation is not only very slow to solve but also numerically unstable. Thus, we decided to scale (and round down) all objective coefficients by a factor of  $10^6$ . This proved to be quite beneficial, with the objective density going from 76 to 45 % and the final dynamism being  $3.627 \times 10^6$ . Still, rounding may change the set of optimal solutions to our problem, in particular if we demand strict tolerances (for example, the optimal objective of 207, 462, 238, 240 is mapped to 207, 392 and not to 207, 462). In order to maintain the correctness of the method regardless of scaling, we implemented the following strategy:

- each subproblem is solved with an objective cutoff of 208, 000.
- feasible solutions are collected but not reported as feasible to the underlying MIP solver, thus forcing the enumeration of all feasible solutions of (scaled) value less than 208, 000, barring symmetry.
- all collected solutions are evaluated with the original objective function and the best one is marked as the optimal solution.

Finally, constraints (10) are not added to the initial formulation, but implemented as indicator constraints, resulting in orders of magnitude savings in the LP solution times.

## 6.3 Symmetry breaking

A preliminary comparison between orbital shrinking and isomorphism pruning showed that the latter is more effective in pruning the search tree, provided that the right branching strategy is used. After some tuning, we found out that a good strategy is to rank the variables  $x_{ijk}$  by decreasing values of  $L_{ijk}$ , using the variable index to break ties. The rationale behind this choice is that in this way we improve the bound quite quickly, because the variables with higher priority are the more expensive ones,

but at the same time those whose contribution is more underestimated by the linear relaxation.

#### 6.4 Cutting planes

We separate local versions of both cutting planes (13) and (25) throughout the tree. However, care must be taken in order to limit the separation overhead, which can be quite substantial (even orders of magnitude in node throughput). We implemented the following strategy:

- the separation procedure for cuts (13) is called at the current node only if the number of variables fixed to 1 is between 2 and 12, and only for those  $x$  variables that are fractional in the LP relaxation
- cuts (13) are added only if a sufficient number [say (10)] of them is violated at the current node. This is to spare the bookkeeping overhead needed by local cuts when only a few can be added. We consider a cut violated if the new coefficient  $L_{ijk}$  is greater than  $1.1w_{ijk}^*/x_{ijk}^*$ , i.e., it must improve upon the current (implied) coefficient by at least 10%.
- the separation procedure for cuts (25) is called at the current node only if the number of variables fixed to 1 is between 2 and 12. In addition, the first variable  $x_{ijk}$  must be fixed to 1, while the second variable  $x_{pqr}$  must be fractional in the LP relaxation. Notice that, given this assumption, we can simplify the cut as

$$w_{ijk} + w_{pqr} \geq T_{ijkpqr}x_{pqr}$$

- also cuts (25) are added only if a sufficient number of them is violated at the current node [again, (10)]. We consider a cut violated if the new coefficient  $T_{ijkpqr}$  is such that

$$T_{ijkpqr}x_{pqr}^* > 1.1(w_{ijk}^* + w_{pqr}^*)$$

- only the LP relaxations of the sub-MIPs needed to compute the coefficients of cuts (13) and (25) are solved, without resorting to enumeration. This has several benefits: (i) an occasionally challenging sub-MIP will not slow down the overall procedure and (ii) the coefficients are weaker and thus the chances of generating a (slightly) invalid cut because of numerics and tolerances are lower.

#### 6.5 Results

We implemented our codes in C++, using IBM ILOG CPLEX 12.5.1 [8] as black box MIP solver through the CPLEX callable library APIs and Perl-lib [20] to implement the group operations needed by isomorphism pruning. All tests have been performed on a cluster of 24 identical machines, each equipped with an Intel Xeon E3-1220v2 running at 3.10 GHz and with 16 GB of RAM.

Computing the symmetry group, as well as enumerating the 85 solutions of model (20)–(24), was done in less than 1 minute (using the constraint programming solver Gecode [4]), and thus can be considered negligible. All sub-MIPs (to be more precise, their LP relaxations) are always solved with default CPLEX tolerances. As far as the subproblem resolution is concerned, we used again default tolerances and disabled all CPLEX cut separators, in order to speedup node throughput (CPLEX cuts were in any case not effective because of symmetry) and improve numerical safety. Aggregated results are reported in Table 4, while detailed results are reported in appendix in Table 3.

According to Table 4, all subproblems required a total of 457,317 s (less than a week) if performed on a serial machine with the same speed. The total number of nodes turned out to be just a little below 2.5 million, surprisingly low given the difficulty of the instance, proving that the techniques developed in the paper were very successful in pruning the enumeration tree. We note that the results refer to the case in which each CPLEX run was given 4 threads and the search mode was set to `CPX_PARALLEL_OPPORTUNISTIC`, resulting in a nondeterministic run. This choice is motivated by the fact that because of our expensive separation procedures, the default load balancing policy of CPLEX resulted in a severe underuse of the available cores. Indeed, forcing deterministic search results in a very similar outcome as far as the number of nodes is concerned, but with a doubled running time.

## 7 Conclusions

We addressed the solution of a very challenging (and previously unsolved) instance of the quadratic 3-dimensional assignment problem, arising in digital wireless communications. In order to solve this instance to optimality, we developed novel techniques, such as a new family of cutting planes and a two-level symmetry handling strategy based on orbital shrinking. These techniques, among others, proved to be instrumental in the successful solution of this Q3AP instance. More importantly, they are general purpose in nature and can be easily extended to different Q3AP or even QAP instances with similar properties.

**Acknowledgments** We thank Peter M. Hahn for providing the numerical data for the 16-PSK Q3AP instance. The work of the first author was supported in part by AFOSR under grant FA9550-12-1-0153.

## 8 Appendix: Detailed computational results

See Table 3.

**Table 3** Subproblem results

Subproblem	Time (s)	Nodes
0,0,8,8,0,0	52,545	599,725
0,1,7,7,1,0	54,943	351,345
0,2,6,6,2,0	45,635	227,453
0,2,6,7,0,1	6,813	22,543
0,3,5,5,3,0	22,055	118,927
0,3,5,6,1,1	13,970	35,419
0,4,4,4,4,0	24,246	133,046
0,4,4,5,2,1	4,150	15,256
0,4,4,6,0,2	89	163
0,5,3,3,5,0	24,977	104,091
0,5,3,4,3,1	1,471	9,025
0,5,3,5,1,2	88	15
0,6,2,2,6,0	18,261	72,035
0,6,2,3,4,1	758	3,099
0,6,2,4,2,2	76	15
0,6,2,5,0,3	20	0
0,7,1,1,7,0	7,398	35,643
0,7,1,2,5,1	245	569
0,7,1,3,3,2	78	15
0,7,1,4,1,3	45	0
0,8,0,0,8,0	452	3,944
0,8,0,1,6,1	103	41
0,8,0,2,4,2	62	15
0,8,0,3,2,3	31	0
0,8,0,4,0,4	14	0
1,0,6,7,2,0	1,843	10,676
1,0,6,8,0,1	818	6,593
1,1,5,6,3,0	6,430	30,187
1,1,5,7,1,1	3,258	12,047
1,2,4,5,4,0	20,217	108,980
1,2,4,6,2,1	1,688	5,940
1,2,4,7,0,2	102	81
1,3,3,4,5,0	33,871	121,652
1,3,3,5,3,1	1,398	6,070
1,3,3,6,1,2	116	17
1,4,2,3,6,0	27,664	87,038
1,4,2,4,4,1	775	2,697
1,4,2,5,2,2	99	15

**Table 3** continued

Subproblem	Time (s)	Nodes
1,4,2,6,0,3	24	0
1,5,1,2,7,0	12,675	46,127
1,5,1,3,5,1	266	433
1,5,1,4,3,2	107	21
1,5,1,5,1,3	47	0
1,6,0,1,8,0	2,464	10,909
1,6,0,2,6,1	154	55
1,6,0,3,4,2	87	19
1,6,0,4,2,3	33	0
1,6,0,5,0,4	17	0
2,0,4,6,4,0	2,494	14,946
2,0,4,7,2,1	226	623
2,0,4,8,0,2	32	25
2,1,3,5,5,0	16,951	92,512
2,1,3,6,3,1	614	1,653
2,1,3,7,1,2	109	15
2,2,2,4,6,0	20,930	94,663
2,2,2,5,4,1	502	1,493
2,2,2,6,2,2	99	15
2,2,2,7,0,3	23	0
2,3,1,3,7,0	11,955	45,794
2,3,1,4,5,1	238	333
2,3,1,5,3,2	101	15
2,3,1,6,1,3	46	0
2,4,0,2,8,0	2,748	10,899
2,4,0,3,6,1	114	43
2,4,0,4,4,2	91	15
2,4,0,5,2,3	34	0
2,4,0,6,0,4	17	0
3,0,2,5,6,0	1,753	10,368
3,0,2,6,4,1	78	49
3,0,2,7,2,2	56	13
3,0,2,8,0,3	11	0
3,1,1,4,7,0	3,627	16,384
3,1,1,5,5,1	111	39
3,1,1,6,3,2	84	13
3,1,1,7,1,3	48	0
3,2,0,3,8,0	1,152	4,732
3,2,0,4,6,1	83	31
3,2,0,5,4,2	69	13
3,2,0,6,2,3	34	0

**Table 3** continued

Subproblem	Time (s)	Nodes
3,2,0,7,0,4	17	0
4,0,0,4,8,0	101	177
4,0,0,5,6,1	43	15
4,0,0,6,4,2	21	0
4,0,0,7,2,3	21	0
4,0,0,8,0,4	6	0
Total	457,312	2,476,819

**Table 4** Aggregated results. Easy subproblems are those requiring less than 100 nodes to solve, while the hard ones are those requiring more than 100,000

Subproblems	Count	Time (s)	Nodes
Easy	43	2,651	610
Medium	34	176,177	710,990
Hard	8	278,489	1,756,219
Total	85	457,317	2,476,819

## References

- Burkard, R., Dell'Amico, M., Martello, S.: Assignment problems. SIAM (2009)
- Fischetti, M., Liberti, L.: Orbital shrinking. In: Mahjoub, A.R., Markakis, V., Milis, I., Paschos, V.T. (eds.) ISCO. Lecture Notes in Computer Science, vol. 7422, pp. 48–58. Springer, New York (2012)
- Fischetti, M., Monaci, M., Salvagnin, D.: Three ideas for the quadratic assignment problem. *Oper Res* **60**(4), 954–964 (2012)
- Gecode Team: Gecode: Generic constraint development environment (2012). Available at <http://www.gecode.org>
- Gent, I.P., Petrie, K.E., Puget, J.F.: Symmetry in constraint programming. In: Rossi, F., van Beek, P., Walsh, T. (eds.) Handbook of Constraint Programming, pp. 329–376. Elsevier, New York (2006)
- Hahn, P.M.: <http://www.seas.upenn.edu/~hahn/>
- Hahn, P.M., Kim, B.J., Stütze, T., Kanthak, S., Hightower, W.L., Samra, H., Ding, Z., Guignard, M.: The quadratic three-dimensional assignment problem: exact and approximate solution methods. *Eur J Oper Res* **184**(2), 416–428 (2008)
- IBM: IBM ILOG Cplex Optimization Studio. <http://www.cplex.com>
- Kaufman, L., Broeckx, F.: An algorithm for the quadratic assignment problem using Benders' decomposition. *Eur J Oper Res* **2**, 204–211 (1978)
- Liberti, L.: Reformulations in mathematical programming: automatic symmetry detection and exploitation. *Math Progr* **131**(1–2), 273–304 (2012)
- Linderöth, J., Margot, F., Thain, G.: Improving bounds on the football pool problem by integer programming and high-throughput computing. *Inf J Comput* **21**(3), 445–457 (2009)
- Margot, F.: Exploiting orbits in symmetric ILP. *Math Progr* **98**(1), 3–21 (2003)
- Margot, F.: Symmetry in integer linear programming. In: Jünger, M., Liebling, T., Naddef, D., Nemhauser, G., Pulleyblank, W., Reinelt, G., Rinaldi, G., Wolsey, L. (eds.) 50 Years of Integer Programming 1958–2008, pp. 647–686. Springer, New York (2010)
- McKay, B.D.: Practical graph isomorphism (1981)
- Mittelman, H.D., Peng, J.: Estimating bounds for quadratic assignment problems associated with hamming and manhattan distance matrices based on semidefinite programming. *SIAM J Optim* **20**(6), 3408–3426 (2010)
- Østergård, P.R.J., Blass, U.: On the size of optimal binary codes of length 9 and covering radius 1. *IEEE Trans Inform Theory* **47**(6), 2556–2557 (2001)

17. Ostrowski, J., Linderoth, J., Rossi, F., Smriglio, S.: Orbital branching. *Math Progr* **126**(1), 147–178 (2011)
18. Peng, J., Mittelman, H.D., Li, X.: A new relaxation framework for quadratic assignment problems based on matrix splitting. *Math Progr Comput* **2**(1), 59–77 (2010)
19. Pierskalla, W.P.: The multi-dimensional assignment problem. *Oper Res* **16**(2), 422–431 (1968)
20. Rehn, T.: Fundamental permutation group algorithms for symmetry computation. Master's thesis, Otto-von-Guericke University Magdeburg (2010)
21. Salvagnin, D.: Orbital shrinking: a new tool for hybrid MIP/CP methods. In: CPAIOR, pp. 204–215 (2013)
22. Salvagnin, D., Walsh, T.: A hybrid MIP/CP approach for multi-activity shift scheduling. In: CP, pp. 633–646 (2012)
23. Samra, H., Ding, Z.: Symbol mapping diversity in iterative decoding/demodulation of ARQ systems. In: ICC, pp. 3585–3589. IEEE (2003)
24. Samra, H., Ding, Z., Hahn, P.M.: Symbol mapping diversity design for multiple packet transmissions. *IEEE Trans Commun* **53**(5), 810–817 (2005)
25. Stützle, T.: Iterated local search for the quadratic assignment problem. *Eur J Oper Res* **174**(3), 1519–1539 (2006)
26. Wu, X., Mittelman, H.D., Wang, X., Wang, J.: On computation of performance bounds of optimal index assignment. *IEEE Trans Commun* **59**(12), 3229–3233 (2011)
27. Xia, Y., Yuan, Y.: A new linearization method for quadratic assignment problem. *Optim Methods Software* **21**, 803–816 (2006)