


# A study on exponential-size neighborhoods for the bin packing problem with conflicts

Renatha Capua<sup>1</sup> · Yuri Frota<sup>1</sup> ·  
Luiz Satoru Ochi<sup>1</sup> · Thibaut Vidal<sup>2</sup> 

Received: 6 April 2017 / Revised: 1 April 2018 / Accepted: 5 April 2018  
© Springer Science+Business Media, LLC, part of Springer Nature 2018

**Abstract** We propose an iterated local search based on several classes of local and large neighborhoods for the bin packing problem with conflicts. This problem, which combines the characteristics of both bin packing and vertex coloring, arises in various application contexts such as logistics and transportation, timetabling, and resource allocation for cloud computing. We introduce  $\mathcal{O}(1)$  evaluation procedures for classical local-search moves, polynomial variants of ejection chains and assignment neighborhoods, an adaptive set covering-based neighborhood, and finally a controlled use of 0-cost moves to further diversify the search. The overall method produces solutions of good quality on the classical benchmark instances and scales very well with an increase of problem size. Extensive computational experiments are conducted to measure the respective contribution of each proposed neighborhood. In particular, the 0-cost moves and the large neighborhood based on set covering contribute very significantly to the search. Several research perspectives are open in relation to possible hybridizations with other state-of-the-art mathematical programming heuristics for this problem.

---

✉ Thibaut Vidal  
vidalt@inf.puc-rio.br

Renatha Capua  
rcapua@ic.uff.br

Yuri Frota  
yuri@ic.uff.br

Luiz Satoru Ochi  
satoru@ic.uff.br

<sup>1</sup> Inst. Computação, Universidade Federal Fluminense, R. Passo da Pátria, 156, Niterói, RJ 24210-240, Brazil

<sup>2</sup> Dep. de Informática, Pontifícia Universidade Católica do Rio de Janeiro, R. Marquês de São Vicente, 225, Rio de Janeiro, RJ 22451-900, Brazil

**Keywords** Metaheuristics · Bin packing with conflicts · Large neighborhood search · Ejection chains · Assignment · Set covering

## 1 Introduction

The Bin Packing Problem with Conflicts (BPPC) is a difficult combinatorial optimization problem which brings together the features of bin packing and vertex coloring problems. We are given a set of bins of identical capacity  $Q$ , a set of items  $\mathcal{I} = \{1, \dots, n\}$  with positive weights  $w_1, \dots, w_n$  and a conflict graph  $G = (V, E)$ . Any edge  $(i, j) \in E$  represents a conflict between items  $i$  and  $j$ . The objective of the BPPC is to assign each item to a bin, in such a way that the number of bins in use is minimized and there are no two items in conflict in the same bin. This problem can be formulated as follows (Gendreau et al. 2004):

$$\min \sum_{k=1}^n y_k \quad (1)$$

$$\text{s.t.} \quad \sum_{i=1}^n w_i x_{ik} \leq Q y_k \quad k = 1, \dots, n \quad (2)$$

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

$$x_{ik} + x_{jk} \leq 1 \quad (i, j) \in E, k = 1, \dots, n \quad (4)$$

$$y_k \in \{0, 1\} \quad k = 1, \dots, n \quad (5)$$

$$x_{ik} \in \{0, 1\} \quad i = 1, \dots, n, k = 1, \dots, n \quad (6)$$

In this model, each binary variable  $x_{ik}$  is set to 1 if and only if item  $i$  is assigned to bin  $k$ , and each binary variable  $y_k$  is set to 1 if and only if the bin  $k$  has been used. Inequality (2) enforces the capacity constraints, Eq. (3) states that each item should be assigned to exactly one bin, and Inequality (4) models the conflict constraints.

The BPPC is linked to several important applications. In the field of transportation and logistics, restrictions and compatibility constraints are imposed on several classes of products, including food, medicine or hazardous materials. Combinations of flammable, explosive, or toxic products should never be carried jointly in the same load (Hamdi-Dhaoui et al. 2014; Minh et al. 2013), leading to difficult item-to-vehicle assignment decisions during operational planning. The problem also arises in examination timetabling, where we search for a schedule that respects room capacities and avoids conflicts (Laporte and Desroches 1984). No student, for example, should take more than one exam at the same time. Finally, another application comes from the field of parallel computing, where we aim to assign a set of tasks to a minimum number of processors, subject to restrictions on the collocation of some conflicting tasks in some machines (Jansen 1999; Masson et al. 2013).

From a methodological standpoint, the problem is also of high interest. It includes two types of constraints: capacity restrictions and conflicts. Each constraint class,

alone, leads to a NP-hard subproblem: bin packing or vertex coloring. It is also noteworthy that the research on the BPPC has been in majority concentrated on approximation algorithms (Epstein and Levin 2008) and mathematical-programming techniques. Mathematical programming has been successful on the BPPC up to this date, since decomposition methods such as branch-and-price (Fernandes Muritiba et al. 2010; Elhedhli et al. 2011; Sadykov and Vanderbeck 2013) are very efficient for problem instances with few items per bin (a common characteristic of most of the test sets available in the literature). Even in these conditions, the current mathematical programming algorithms remain applicable only to medium-scale problems, and many instances with one thousand items cannot be exactly solved to this date. This is unfortunately insufficient for several applications of BPPC, for cloud computing and large-scale task allocation problems (Masson et al. 2013), which can involve tens or hundreds of thousands of items.

The objective of this article is to progress on the understanding and development of metaheuristics for this class of problems. A significant part of recent research on metaheuristics, in the last 5 years, has been focused on producing *novel* search concepts, often from natural analogies. However, as in Sörensen (2015), we believe that part of this research is misled. Beyond natural phenomena, the success of many metaheuristics come from a careful definition of their neighborhoods, efficient techniques to *search* them, and an exploration strategy which finds a sensible balance between intensification and diversification (Blum and Roli 2003; Vidal et al. 2013). Metaheuristics are also better understood in terms of their basic operations: neighborhoods, local search, restarts and perturbations, recombinations, short and long-term memories and diversification procedures, among others. Following this general logic, we develop a simple metaheuristic framework, an iterated local search, and enhance this search procedure with a variety of local and large (exponential-sized) neighborhoods, with the aim of carefully investigating their relative contribution to the search. We also propose efficient evaluation procedures in  $\mathcal{O}(1)$  for standard enumerative neighborhoods such as SWAP or RELOCATE, therefore avoiding iterating over conflicting items in a time proportional to the degree of the vertices. Overall, this leads to a simple and efficient algorithm, plus a more complete version with additional classes of neighborhoods, which both return results of good quality on classic instances for the considered problem. We finally conduct extensive computational analyses in order to measure the contribution of each neighborhood in the search and its relative computational time consumption. In particular, we observe that set covering-based neighborhoods significantly contribute to the search performance, as well as diversification strategies based on  $\theta$ -cost moves.

Overall, the main contributions of this paper are:

1. A simple and efficient ILS approach for the BPPC;
2. Sophisticated adaptations of exponential-size neighborhoods (ejection chains, assignment, and grenade) for the problem;
3. Efficient amortized  $\mathcal{O}(1)$  move evaluations for standard enumerative neighborhoods;
4. A controlled use of  $\theta$ -cost moves to better explore the search place;
5. An investigation of the relative contribution of all proposed neighborhoods.

## 2 Related literature

Although conflict constraints have a long history in the domain of timetabling (Laporte and Desroches 1984), the first studies dedicated to the bin packing problem with conflicts appeared in Jansen and Öhring (1997) and Jansen (1999). Early works were focused on approximation algorithms, a research line which was later pursued in Epstein et al. (2008) and Epstein and Levin (2008). Unfortunately, no bounded polynomial approximation scheme can be obtained for the BPPC with general conflict graphs (Jansen 1999). As such, most approximation results concern specific graph structures, such as perfect graphs, interval graphs, and bipartite graphs, with approximation ratios of 2.5, 2.33333, and 1.75, respectively (Epstein and Levin 2008).

More recently, Gendreau et al. (2004) investigated some heuristics (without performance guarantee) for the problem. The authors proposed six different constructive heuristics, as well as two lower bounds, and a first set of benchmark instances based on the data set of Falkenauer (1996) for the bin packing problem. Unfortunately, these instances are not available anymore to this date, and thus were re-generated later on (Fernandes Muritiba et al. 2010; Elhedhli et al. 2011).

Subsequently, a significant research effort has been dedicated to mathematical programming algorithms for the problem. Khanafer et al. (2010) proposed new lower bounds based on the concept of data-dependent dual feasible solutions. Furthermore, three recent articles proposed exact branch-and-price procedures (Fernandes Muritiba et al. 2010; Elhedhli et al. 2011; Sadykov and Vanderbeck 2013). These articles use different branching rules, initial columns, and sub-procedures for the pricing problem (a knapsack problem with conflicts). When the conflicts form an interval graph, the pricing problem can be solved efficiently via dynamic programming (Sadykov and Vanderbeck 2013), while branch-and-bound procedures can be used in more general cases. These exact methods can solve many benchmark instances, with up to 1000 items in some cases. Still, instances with a large number of items per bin (e.g., set (da) of Sadykov and Vanderbeck 2013) remain challenging. The computational effort can also vary significantly from one instance to another, and becomes impracticable for large data sets with general graphs.

Metaheuristics for the BPPC have received, to this date, less attention than mathematical-programming algorithms. An advanced population-based metaheuristic was proposed in Fernandes Muritiba et al. (2010), and used to generate good upper bounds and initial columns. The method is a complex combination of a tabu search based on the impasse class neighborhoods of Morgenstern (1996), with a genetic algorithm using the crossover operator of Galinier and Hao (1999). Sadykov and Vanderbeck (2013) also introduce a diving heuristic based on a controlled partial exploration of the branch-and-price search tree. This method is a good alternative between exact solution approaches and heuristics, as it combines the benefits of mathematical programming for problem instances with few items per bin, with a shorter CPU time. Overall, metaheuristics for the BPPC still deserve further investigation.

Finally, as the BPPC is a generalization of both the bin packing problem and vertex coloring, results about metaheuristic searches for these two subproblems can be a good source of inspiration. The literature on these problems is extensive, and we refer to Lewis et al. (2012), Malaguti and Toth (2010), Delorme et al. (2016), Quiroz-

Castellanos et al. (2015) and Lewis (2016) for some detailed reviews. Some of the current best metaheuristics for the bin packing problem include the grouping genetic algorithm of Falkenauer (1996), the extension of the Perturbation-MBS method by Fleszar and Charalambous (2011), and the grouping genetic algorithm with controlled gene transmission of Quiroz-Castellanos et al. (2015). Population-based heuristics tend to be frequently used, due to their capability of exploring diverse areas of the search space. For the vertex coloring problem, some of the state-of-the-art metaheuristics include the TabuCol approach of Hertz and de Werra (1987), the PartialCol algorithm of Blöchliger and Zufferey (2008) with fluctuation of the objective function, the hybrid evolutionary algorithms of Galinier and Hao (1999), Malaguti et al. (2008), and the ant colony optimization metaheuristic of Dowsland and Thompson (2008).

### 3 Methodology

We opted for a simple metaheuristic framework, which allows to carefully investigate the impact of several local and large neighborhoods. Our method is thus built on an Iterated Local Search (ILS) metaheuristic, with several neighborhood classes for solution improvement, and one perturbation procedure. Its pseudo-code is presented in Algorithm 1. Note that many methodological improvements concerning local-search methods are generalizable to other metaheuristics, such that an extension of the proposed approach within a variable neighborhood search (considering multiple perturbation procedures), or its integration into a hybrid genetic algorithm (using a local search for solution improvement during mutations) constitute promising open research avenues.

An initial solution is first generated via a constructive heuristic (Line 1). Subsequently, to decrease the number of used bins, the method iteratively removes one bin (Line 4), randomly assigning its items to other bins, and attempts to solve the resulting conflicts or capacity excesses by means of iterations of local search with 0–cost moves, ejection chains, grenade moves, resolutions of set covering or assignment formulations, and perturbation procedures (Lines 7–17). Throughout the search, the solution with the least penalty ( $\mathcal{S}_{\text{BEST}}$ ) is recorded. Using the terminology of the vertex coloring problem, and as surveyed in Lewis et al. (2012), this type of strategy exploits the search space of *complete, but improper* colorings. The algorithm terminates when no feasible solution is found after  $N_{\text{SHAK}}$  perturbations, or if a trivial lower bound  $K_{\text{LB}}$  on the number of bins (the sum of items weight divided by bin capacity) is reached. We now describe each component of the algorithm in further details.

#### 3.1 Initial solution

The initial solution is generated by means of the modified first fit decreasing heuristic (Gendreau et al. 2004). Items are enumerated in non-decreasing order of weight, and inserted in the first bin which has enough residual space and does not contain conflicting items. This leads to a feasible initial solution with  $K$  bins.

**Algorithm 1** Iterated local search for the BPPC

---

```

1:  $\mathcal{S}_{\text{BEST}} \leftarrow \text{BuildInitialSolution}();$ 
2:  $K_{\text{LB}} \leftarrow \text{ComputeLowerBound}();$ 
3: while  $\text{NbBins}(\mathcal{S}_{\text{BEST}}) > K_{\text{LB}}$  and  $\mathcal{S}_{\text{BEST}}$  is feasible do
4:    $\mathcal{S}_{\text{BEST}} \leftarrow \text{ReduceNbBins}(\mathcal{S}_{\text{BEST}});$ 
5:    $\mathcal{S} \leftarrow \mathcal{S}_{\text{BEST}};$ 
6:    $i_{\text{SHAK}} \leftarrow 0;$ 
7:   while  $i_{\text{SHAK}} < N_{\text{SHAK}}$  and  $\mathcal{S}_{\text{BEST}}$  is not feasible do
8:     for  $N_{\text{LS}}$  iterations do
9:        $\mathcal{S} \leftarrow \text{LocalSearch}(\mathcal{S});$ 
10:       $\mathcal{S} \leftarrow \text{AssignmentNeighborhood}(\mathcal{S});$ 
11:       $\mathcal{S} \leftarrow \text{EjectionChains}(\mathcal{S});$ 
12:       $\mathcal{S} \leftarrow \text{Grenade}(\mathcal{S});$ 
13:     end for
14:     if  $\exists k \in \mathbb{N}^+$  such that  $i_{\text{SHAK}} = k \times N_{\text{SC}}$  then  $\mathcal{S} \leftarrow \text{SetCovering}();$ 
15:     if  $\text{cost}(\mathcal{S}) < \text{cost}(\mathcal{S}_{\text{BEST}})$ , then  $\mathcal{S}_{\text{BEST}} \leftarrow \mathcal{S}$  and  $i_{\text{SHAK}} \leftarrow 0;$ 
16:      $\mathcal{S} \leftarrow \text{Shaking}(\mathcal{S}_{\text{BEST}});$ 
17:      $i_{\text{SHAK}} \leftarrow i_{\text{SHAK}} + 1;$ 
18:   end while
19: end while

```

---

### 3.2 Local search

The local search procedure aims to improve any infeasible solution produced by the removal of one bin or by the shaking operator. To define an improvement, the cost of each solution  $\mathcal{S}$  is evaluated as the sum of the cost of its bins  $B \in \mathcal{S}$  given by Eq. (7). In this equation,  $C(B)$  represents the current number of conflicts in the bin,  $W(B)$  represents the total excess weight, and  $\omega^C$  and  $\omega^W$  are the penalty factors associated with each unit conflict and weight violation.

$$\Phi(B) = \omega^C C(B) + \omega^W W(B) \quad (7)$$

*Neighborhoods and exploration.* Three classic neighborhoods are used: SWAP, RELOCATE and SWAP2VS1. A SWAP move exchanges two items from different bins. As suggested by its name, a RELOCATE move reassigns an item to a different bin. Finally, a SWAP2VS1 move exchanges a pair of consecutive items from one bin with a single item from another bin. The neighborhoods are explored as detailed in Algorithm 2, enumerating the pairs of bins in random order to test the associated moves (Lines 4–5). For each bin pair, the best improving move is applied if such a move exists (Line 9), and the local search terminates when no improving move can be found.

This evaluation policy allows to save computational effort, since it is unnecessary to re-evaluate the moves between a pair of bins when the evaluation has been done in the past without success, and if the bins have not undergone any change. Moreover, at least one bin  $B$  involved in a move must satisfy  $\Phi(B) > 0$  to hope for an improvement.

This condition is included as an additional filter (Line 4), allowing to save a significant amount of CPU time. Finally, during the first loop of the local search, the method accepts  $0$ -cost moves (Line 8). These moves do not improve the objective but contribute to diversify the search.

---

**Algorithm 2** Local search with 0-cost moves

---

```

1:  $It_{LOOP} \leftarrow 0$ 
2: while a local minimum has not been attained do
3:    $It_{LOOP} \leftarrow It_{LOOP} + 1$ 
4:   for each bin  $B$  in random order such that  $\Phi(B) > 0$  do
5:     for each bin  $B' \neq B$  in random order do
6:       if  $B$  or  $B'$  has been modified since the last evaluation or  $It_{LOOP} = 1$  then
7:          $\Delta \leftarrow \text{CostBestMove}(B, B')$ 
8:         if  $\{\Delta < 0\} \text{or} \{\Delta \leq 0 \text{ and } It_{LOOP} = 1\}$  then
9:            $\text{ApplyBestMove}(B, B')$ 
10:        end if
11:       end if
12:     end for
13:   end for
14: end while

```

---


*Efficient move evaluations.* The computational efficiency of the move evaluations is of critical importance, as it is the main bottleneck of most recent local search-based heuristics. Any significant reduction in the CPU time needed to evaluate moves can be translated into a direct gain in terms of number of moves which can be evaluated during one run of the algorithm, and thus into enhanced solution quality. For the BPPC, move evaluations are not trivial since they require to compute the number of conflicts in the new solution. A straightforward implementation of SWAP, for example, between bins  $B$  and  $B'$  would require  $\mathcal{O}(|B| + |B'|)$  elementary operations per move evaluation, where  $|B|$  is the current number of items in a bin  $B$ . This effort is due to the evaluation of conflicts between the items exchanged and those existing in the bins. This effort grows linearly with the average number of items per bin, and can be responsible for the majority of the CPU time for problem instances with a large number of items per bin.

To address this issue, we propose a pre-processing and move evaluation mechanism which supports amortized  $\mathcal{O}(1)$  move evaluations. We rely on an array which gives  $\mathcal{O}(1)$  access to the current number of conflicts  $\text{CONF}[i][B]$  between an item  $i$  and all items in bin  $B$ . Let  $d_i$  be the degree of  $i$  in the conflict graph. This array can be built in  $\mathcal{O}(Kn + \sum_{i=1}^n d_i)$  elementary operations for the initial solution. Subsequently, any solution modification can be assimilated to a sequence of item relocations. For each relocation of an item  $i$  from a bin  $B$  to a bin  $B'$ , the CONF array is updated in  $\mathcal{O}(d_i)$  by means of Algorithm 3.

This data structure can now be used to perform efficient move evaluations. Consider a SWAP move between an item  $i$  from bin  $B$ , with item  $j$  from bin  $B'$ . The difference

**Algorithm 3** Update mechanism – Relocation of one item  $i$  from bin  $B$  to  $B'$

- 1: **for** each item  $j$  in conflict with  $i$  **do**
- 2:    $\text{CONF}[j][B] = \text{CONF}[j][B] - 1$
- 3:    $\text{CONF}[j][B'] = \text{CONF}[j][B'] + 1$
- 4: **end for**



**Fig. 1** A generalized move involving the exchange of three items in  $B$  with three items in  $B'$

$\Delta$  in the number of conflicts after and before the move can be computed as follows:

$$\Delta = \text{CONF}[i][B'] + \text{CONF}[j][B] - \text{CONF}[i][B] - \text{CONF}[j][B'] - 2 \times \text{ISCONFLICT}(i, j) \tag{8}$$

Remark the corrective term  $2 \times \text{ISCONFLICT}(i, j)$ , which takes value 2 if and only if there is a conflict between the items  $i$  and  $j$ . This term comes from the fact that the array  $\text{CONF}[i][B']$  contains the number of conflicts when moving  $i$  in  $B'$ , but *before* the removal of item  $j$ , and vice-versa for the array  $\text{CONF}[j][B]$ .

This approach can be generalized to evaluate in  $\mathcal{O}(1)$  any move involving an exchange of bounded sets of items  $S$  and  $S'$  between a pair of bins  $B$  and  $B'$ , hence allowing efficient evaluations for all the considered moves. An example of a move, which exchanges three items of each bin, is displayed in Fig. 1. Let  $N_{\text{INTER}}(S, S')$  be the number of conflicts between pairs of items in  $S$  and  $S'$ , let  $N_{\text{INTRA}}(S)$  be the number of conflicts within the set  $S$ , and let  $N_{\text{INTRA}}(S')$  be the number of conflicts within the set  $S'$ . The difference in the number of conflicts after the application of the move can be obtained from Eq. (9).

$$\Delta = \sum_{i \in S} (\text{CONF}[i][B'] - \text{CONF}[i][B]) + \sum_{j \in S'} (\text{CONF}[j][B] - \text{CONF}[j][B']) + 2 \times (N_{\text{INTRA}}(S) + N_{\text{INTRA}}(S') - N_{\text{INTER}}(S, S')) \tag{9}$$

Capacity excesses can also be simply evaluated in  $\mathcal{O}(1)$ . Overall, this leads to move evaluations in  $\mathcal{O}(1)$  at the price of a slightly higher computational effort when updating the solution. In all computational experiments, we observed that hundreds of moves



are tested, in average, before even applying one of them. The complexity needed to initialize and update the data structures remained largely dominated by the effort dedicated to move evaluations, leading to the announced amortized  $\mathcal{O}(1)$  evaluations.

### 3.3 Large neighborhoods


Our hybrid ILS is built on simple perturbation and local search concepts, but it relies on four large neighborhoods to enhance the search performance. The first two neighborhoods—ASSIGNMENT and EJECTION CHAINS—exploit restrictions of the problem which have the benefit to be combinatorial but still polynomial. This leads to neighborhoods of exponential size which can be searched in polynomial time. The third neighborhood—GRENADE—is based on the enumeration of a larger move. Finally, the fourth neighborhood—SET COVERING—uses the fact that effective integer programming solvers can address the set covering formulation of the problem in the presence of a limited pool of columns (item combinations in a bin). These solvers are used to combine high-quality columns from local minima to produce better solutions.

Some of these families of neighborhoods are well known in the operations research literature, and we refer to Deineko and Woeginger (2000) and Ahuja et al. (2002) for detailed surveys and analyses. Not all combinatorial optimization problems allow for an efficient search of such exponential-sized neighborhoods. For the case of the BPPC, such a search is possible. We now describe these large neighborhoods, tailored for the problem, as well as a specific form of ejection chains, which also allows a polynomial exploration of an exponential subset of solutions, under the condition that ejections between bins should follow a pre-specified ordering.

#### 3.3.1 Assignment neighborhood

The assignment neighborhood selects some items for removal, and reinserts these items in the *holes* thus formed by solving an assignment problem in a weighted bipartite graph. Such a methodology was successfully used for traveling salesman and vehicle routing problems in Sarvanov and Doroshko (1981), Gutin (1999) and Toth and Tramontani (2008).

We adapt this approach to the BPPC as follows. In our context, only a few *problematic* items usually present conflicts or belong to a bin with excess weight during the course of the search. To improve the solution, better allocation decisions should be sought for some of these items with simultaneous movements of other items. Hence, the proposed algorithm randomly selects one problematic item, call it  $k$ , as well as  $N_{\text{ASSIGN}}$  random items, at most one in each bin. Together, these items form a vertex set  $V'_1$ . Let  $V'_2$  be a copy of this set of vertices, and define the edge set  $E' = \{(i, j) \text{ such that } i \in V'_1 \text{ and } j \in V'_2\}$ . The cost of each edge  $(i, j)$  is defined as the cost of inserting item  $i$  in the bin  $B(j)$  of item  $j$  (which has been removed beforehand). Finally, to break possible ties and favor a movement of  $k$ , an  $\epsilon$ -penalty is added to the cost of the edge linking  $k$  with its copy. This defines a bipartite graph  $G' = (V'_1, V'_2, E')$ , illustrated in Fig. 2.



**Fig. 2** Assignment neighborhood and its associated bipartite graph  $G'$

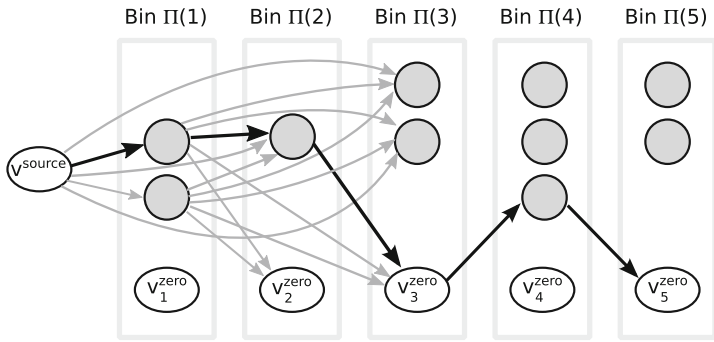
Then, finding a best re-allocation of items to bins is done by solving a linear-sum assignment problem on this graph. The optimal solution is found in  $\mathcal{O}(n^3)$  by means of an efficient implementation of the Hungarian algorithm (Kuhn 1955). Any non-trivial optimal solution of this assignment model (different from the pairing of each node  $i$  with its clone) corresponds to a move of the BPPC solution. This move is applied, leading to a new (better or equal) incumbent solution in the ILS.

### 3.3.2 Ejection chains

Ejection chains allow to explore a different subset of neighbor solutions obtained via chained relocations of items (see, e.g., Thompson and Psaraftis 1993, Glover 1996 and Deineko and Woeginger 2000). A key difference with the previous assignment neighborhood is that the choice of items to be relocated is not a-priori fixed. On the other hand, to allow for a polynomial exploration procedure, we restrain item relocations so as to comply with a pre-defined bin ordering.

Our ejection chains algorithm works as follows. First, choose a random order  $\Pi$  of the bins in the current solution. Then, define an auxiliary graph  $G'' = (V'', A'')$  with  $V'' = \mathcal{V} \cup \mathcal{V}^{\text{ZERO}} \cup v^{\text{SOURCE}}$ , illustrated in Fig. 3.

The set  $\mathcal{V}$  contains one node for each item, while  $\mathcal{V}^{\text{ZERO}}$  contains one node  $v_k^{\text{ZERO}}$  for each bin  $k$  in the current solution. Each node in  $\mathcal{V}$  will model the possible replacement of one item by another, and each node in  $\mathcal{V}^{\text{ZERO}}$  will model the possible insertion of an item and the end of an ejection chain. The node  $v^{\text{SOURCE}}$  represents a source. We denote by  $B(i)$  the bin that contains vertex  $i$  in the current solution. The set  $A''$  contains one arc  $(i, j)$  for any pair of nodes  $i \in \mathcal{V} \cup \mathcal{V}^{\text{ZERO}}$  and  $j \in \mathcal{V} \cup \mathcal{V}^{\text{ZERO}}$  such that bin  $B(i)$  precedes bin  $B(j)$  in the order  $\Pi$ , and one arc  $(v^{\text{SOURCE}}, j)$  for any  $j \in \mathcal{V}$ . The costs of the arcs are defined as follows:



**Fig. 3** Ejection-chains graph, and a possible solution

$$c_{ij} = \begin{cases} \text{the cost difference of bin } B(j) \text{ when replacing item } j \text{ by item } i & \text{if } i \in \mathcal{V} \text{ and } j \in \mathcal{V}, \\ \text{the cost difference of bin } B(j) \text{ when removing item } j & \text{if } i \in \mathcal{V}^{\text{ZERO}} \cup v^{\text{SOURCE}} \text{ and } j \in \mathcal{V}, \\ \text{the cost difference of bin } B(j) \text{ when inserting item } i & \text{if } i \in \mathcal{V} \text{ and } j \in \mathcal{V}^{\text{ZERO}}. \end{cases}$$

A shortest path in this graph from the source to any node in  $\mathcal{V}^{\text{ZERO}}$  is equivalent to a sequence of combined item insertions and removals. Note that, thanks to the use of the nodes in  $\mathcal{V}^{\text{ZERO}}$ , we can possibly obtain a collection of several disjoint ejection chains. The solution is finally updated in case of improvement.

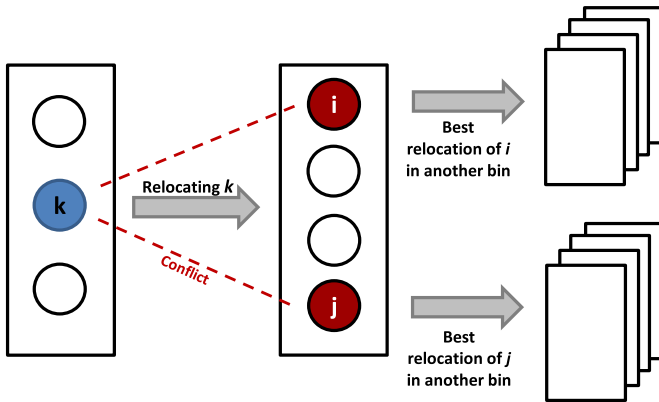
### 3.3.3 Grenade neighborhood

The GRENADE neighborhood has been previously used for the vertex coloring problem in Avanthay et al. (2003), and consists of the enumeration of some larger moves. This neighborhood was selected due to its simplicity and low evaluation complexity. In our context, we consider in turn each item  $k$  in the set of problematic items  $\mathcal{P}$  (as defined in the previous section: with either a conflict or in a bin with excess capacity). For each possible bin  $B \neq B(k)$ , the method evaluates the possibility of relocating  $k$  in  $B$  and jointly relocating any conflicting item  $i$  of  $B$  to another bin by means of a best insertion criterion. The total cost of these combined relocations is evaluated and the best grenade move for  $k$  is applied in case of improvement.

The overall scheme is illustrated in Fig. 4, and the pseudo-code of the neighborhood exploration is detailed in Algorithm 4. Note that the search consists of a single loop over the items  $k \in \mathcal{P}$ , as a complete descent towards a local minimum of this neighborhood would be time consuming.

### 3.3.4 Adaptive set covering

Set covering (or set partitioning) formulations have been successfully used to solve a variety of combinatorial optimization problems to optimality and to generate larger neighborhoods (see, e.g., Monaci and Toth 2006; Muter et al. 2010 and Subramanian et al. 2013 for recent applications on a wide range of packing, coloring and routing problems). In our context, we investigate a set covering formulation in which the



**Fig. 4** GRENADE move. Relocating the item  $i$  creates two conflicts which are resolved by other relocations

---

**Algorithm 4** Exploration of the GRENADE neighborhood

---

```

1: for each item  $k \in \mathcal{P}$  in random order do
2:   for each bin  $B \neq B(k)$  do
3:      $\Delta \leftarrow$  Cost of a relocation of item  $k$  in  $B$ 
4:     for each item  $i \in B$  in conflict with  $k$ , in random order do
5:        $\Delta \leftarrow \Delta +$  Least-cost relocation of  $i$  in a bin different from  $B$ 
6:     end for
7:     UpdateBestGrenadeMove( $\Delta, k$ )
8:   end for
9:   if BestGrenadeMove( $k$ ) < 0, then apply this move
10: end for

```

---

maximum number of bins is fixed, and one aims to minimize the sum of column costs associated to possible capacity of conflict constraints violations, and use it as a base for an additional large neighborhood in the ILS. To limit the computational effort, this procedure is used only a few times during the search, after each  $N_{sc}$  consecutive iterations of the local search procedure without improvement.

The considered set covering formulation is expressed in Eqs. (10–13). It involves the set  $R$  of all possible columns, where each column represents a combination of items, associated with a cost  $c_j$ . The cost  $c_j$  is based on the same weighted penalties for conflicts and capacity excess as in Eq. (7), and let  $R_i \subseteq R$  be the subset of columns containing an item  $i \in \mathcal{V}$ . The goal of the formulation is to find a set of up to  $K$  columns that covers all items with minimum cost. We highlight the fact that the positive-cost columns (with  $c_i > 0$ ) correspond to sets of items with excess capacity or conflicts, and that any solution with cost 0 would correspond to a feasible packing in  $K$  bins.

$$\min \sum_{j \in R} c_j u_j \quad (10)$$

$$\text{s.t. } \sum_{j \in R_i} u_j \geq 1 \quad \forall i \in \mathcal{V} \quad (11)$$

$$\sum_{j \in R} u_j \leq K \quad (12)$$

$$u_j \in \{0, 1\} \quad \forall j \in R \quad (13)$$

This formulation involves an exponential number of decision variables ( $2^n - 1$ ). To circumvent this issue, we restrict the model to a subset of columns  $R' \subset R$ , obtained from local minima of the ILS. More specifically, before every shaking iteration, the columns of the current solution are added to the pool  $R'$ , filtering possible duplicates (this can easily be implemented by hashing), and controlling the size of the pool to remain below a limit  $S_{\text{POOL}}$ . The new columns start to replace the oldest columns when this limit is attained. The formulation is solved by means of an integer programming solver subject to a CPU time limit  $T_{\text{LIMIT}}$ . The pool-size limit is adapted during the search, with the aim to balance the difficulty of the problem and thus obtain good solutions with a formulation that remains tractable. If the model was solved to the optimality during the allowed time, then  $S_{\text{POOL}}$  is increased by 15%. On the contrary, if the time limit was attained without an optimal solution, then  $S_{\text{POOL}}$  is reduced by 15%. The optimal solution provided by this model is taken as new incumbent solution.

### 3.4 Perturbation

Shaking is only operated after  $N_{\text{LS}}$  consecutive iterations of the 0-cost moves, local and large neighborhood search without improvement of the best solution. Indeed, the 0-cost moves themselves already allow significant modifications to the solutions, which help to progress towards different areas of the search space.

During shaking,  $S_{\text{SHAK}}$  items are randomly relocated, considering 50% items with conflicts or belonging to a bin with capacity excess, and 50% random items. The size of the shaking operator is a key parameter for the search. Its value has been defined as a result of our preliminary computational experiments, as described in the next section.

## 4 Computational experiments

The goal of these experiments is to evaluate the performance of the proposed method on the BPPC, as well as to identify the relative contribution of each neighborhood and the impact of its key design choices to drive future research. The classical benchmark instances for the BPPC are divided into six sets: (t) and (u) from Fernandes Muritiba et al. (2010), and (ta), (ua), (d), (da) from Sadykov and Vanderbeck (2013), for a total of 2060 instances ranging from 60 to 1000 items.

The instances of the group (t – triplet) and (u – uniform) were first derived from the bin packing instances of Falkenauer (1996) by Gendreau et al. (2004). The generation procedure, however, contains some randomized choices and the original instances are no longer available. Thus, these instances were generated again in Fernandes Muritiba et al. (2010) and Elhedhli et al. (2011). In this paper, we rely on the files of Fernandes

Muritiba et al. (2010), as this allows a comparison with the only existing population-based metaheuristic in the literature, as well as the current state-of-the-art mathematical programming—heuristic and exact—algorithms of Sadykov and Vanderbeck (2013). The set (t) includes four groups of instances with a number of items ranging from 60 to 501. One remarkable characteristic of this instance group is that, in the optimal bin packing solution, without conflict constraints, the bins are filled by exactly three items each. In a similar manner, the set (u) includes four classes of instances, with 120–1000 items. In both groups, the conflict graph is an interval graph, with a density ranging from 0% to 90%. Previous articles did not always report results for the instances with 0% density, leading to 360 instances per set (9 density levels  $\times$  4 size levels  $\times$  10 instances).

The instances of sets (ta) and (ua) have the same characteristics as the instances of sets (t) and (u), but with an arbitrary conflict graph. These instances were generated by Sadykov and Vanderbeck (2013) with edge densities from 10 to 90%, for a total of 360 instances in each set. Finally, the instances of sets (d) and (da) were generated with item weights uniformly distributed in [500, 2500], and a bin capacity of 10,000. These instances include either 120, 250 or 500 items, with interval (d) or arbitrary (da) conflict graph. The objective of these instances was to include a higher average number of items per bin and obtain more difficult problems for column generation-based methods. This objective was, however, only partially attained, since the average number of items per bin in the best known solutions remain relatively small for all data sets: 2.18, 1.96, 2.91, 2.48, 2.76, and 5.83 for the classes (t), (u), (ta), (ua), (d), and (da), respectively. For this reason, these instance sets remain ideal for existing branch-and-price methods.

#### 4.1 Performance evaluations

The first experiment aims to measure the performance of the proposed ILS with large neighborhoods relatively to the performance of the best current metaheuristic, the population heuristic (PH) of Fernandes Muritiba et al. (2010), and the best mathematical-programming based heuristics: the diving approach (DH) of Sadykov and Vanderbeck (2013) and its variant with limited discrepancy search (DH-LDS). As the total number of instances for the BPPC is very large, the results are presented in an aggregated form. For a precise comparison, the detailed results of PH, DH and DH-DLS were kindly provided by the authors in a personal communication. The proposed algorithm was coded in C++, uses CPLEX 12.5.1, and was run on a single thread of a Xeon X5675 3.07 GHz CPU. PH was run on a Pentium IV 3.0GHz. Finally, DH and DH-LDS were run on a Xeon X5460 3.16GHz. To account for speed differences between processors, all time measures for PH, DH and DH-DLS in this article are multiplied by a scaling factor based on Passmark single-thread benchmark:  $\times 0.463$  for PH, and  $\times 0.952$  for DH and DH-DLS.

A preliminary calibration of the search parameters of the ILS—size and frequency of the shaking operator, number of iterations of the method, size of the set covering subproblem—was conducted with the aim of producing good solutions in a CPU time which is comparable with previous research for medium-sized instances. For this purpose, we varied each main parameter, in turn, to attain a *local minimum* in

terms of solution quality. This led to a standard parameter configuration:  $N_{\text{SHAK}} = 50$ ,  $N_{\text{LS}} = 100$ ,  $N_{\text{SC}} = 25$ ,  $S_{\text{POOL}} = 1500$ ,  $T_{\text{LIMIT}} = 20$  and  $S_{\text{SHAK}} = 3$ . The impact of any deviation from these values will be investigated in deeper details in Sect. 4.2.

To increase the statistical strength of the observations, the proposed method was run ten times for each instance, with different random seeds, in order to measure its average performance. Therefore, all measures of the ILS (solution quality, time, number of optimal solutions found) are averaged over the ten runs. In contrast, previous algorithms (PH, DH and DH-LDS) were run once for each instance. The complete version of the algorithm is called *HILS-Complete*. To examine the impact of the large neighborhood search on solution quality, we will also report the results of a simplified configuration, *ILS-Simple*, which corresponds to the same algorithm without the large neighborhoods. These experiments are presented in Tables 1 and 2. In Table 1, the results are aggregated for each instance class  $\times$  number of items, while Table 2 aggregates the results per instance class  $\times$  density. The following statistics are reported:

**Gap**—The average gap (per instance and run) from the best known solutions (BKS) in the literature. For each instance and run, the gap is calculated as  $100(z - z_{\text{BKS}})/z_{\text{BKS}}$ , where  $z$  is the value of the solution found by the method and  $z_{\text{BKS}}$  is the value of the BKS;

**T(s)**—The average CPU time (per instance and run) until termination of the search;

**Opt**—The average number of optimal solutions (per run) found for this group of instances.

For each group of instances, the best results are highlighted in boldface. The last two columns indicate the average number of bins in the best known solutions of the literature, and the number of known optimal solutions for the group. Finally, the last two rows report general statistics over the sets of instances (t) + (u), and all sets except (d), respectively.

A number of factors are determining for the performance of an optimization approach: the quality of its solutions, its average computational effort on a determined set of instances, and the scalability of the method, i.e., how fast the CPU time grows for larger problem instances. Based on our experimental analyses, the following observations can be made.

In terms of solution quality, *HILS-Complete* provides solutions of higher quality than the current best metaheuristic (PH) and the column generation-based diving algorithm (DH), and slightly lower quality than the diving algorithm with limited discrepancy search (DH-LDS). The magnitude of the differences depends on the sets of instances. For the sets (t) and (u), which are common to all algorithms, *HILS-Complete* finds an average gap of 0.02%, compared to 0.29, 0.18 and 0.00 for PH, DH and DH-LDS, respectively. For the sets (ta) and (ua), the difference of solution quality is also small, and both *HILS-Complete* and DH-LDS return near-optimal solutions with average gaps below 0.15%. For the (da) instances, DH-LDS produces solutions of higher quality than other methods, albeit at the expense of a larger CPU time. Overall, *ILS-Simple*, DH, *HILS-Complete*, DH-LDS produce non-dominated trade-offs in terms of computational effort and solution quality. *HILS-Complete* attained 1819 out of the 1917 known optimal solutions, and a total of 9 new best upper bounds were also found for some open (da) instances.

**Table 1** Comparison of the current state-of-the-art BPPC heuristics for each instance set and size  $n$

Instance	PH			DH			DH-LDS			ILS-Simple			HILS-Complete			BKS								
	Class	$n$	Gap	T(s)	Opt	Gap	T(s)	Opt	Gap	T(s)	Opt	Gap	T(s)	Opt	Gap	T(s)	Opt	Gap	T(s)	Opt	Bin	Opt		
t	60	0.50	19.31	81	0.18	0.18	0.19	90	0.42	0.72	82.1	<b>0.00</b>	1.26	90.0	<b>0.00</b>	1.26	90.0	33.40	90	33.40	90	90.0	90	
	120	0.66	20.59	66	0.48	0.88	1.29	89	0.63	2.34	66.3	<b>0.01</b>	5.17	89.6	<b>0.01</b>	5.17	89.6	66.11	90	66.11	90	89.6	90	
	249	0.44	26.69	58	0.29	5.40	7.43	90	0.35	5.09	57.5	<b>0.08</b>	17.27	84.3	<b>0.08</b>	17.27	84.3	135.83	90	135.83	90	84.3	90	
	501	0.23	30.34	60	0.16	33.06	47.99	90	0.26	13.19	47.9	<b>0.01</b>	23.88	88.9	<b>0.01</b>	23.88	88.9	275.69	90	275.69	90	88.9	90	
u	120	0.11	11.39	85	0.16	0.59	0.72	90	0.06	1.22	87.1	<b>0.03</b>	3.67	88.8	<b>0.03</b>	3.67	88.8	70.38	90	70.38	90	88.8	90	
	250	0.19	23.08	71	0.10	2.98	3.18	90	0.13	4.67	74.9	<b>0.01</b>	17.57	88.8	<b>0.01</b>	17.57	88.8	143.72	90	143.72	90	88.8	90	
	500	0.18	30.78	65	0.05	16.79	17.83	90	0.37	16.47	45.8	<b>0.01</b>	37.56	86.8	<b>0.01</b>	37.56	86.8	286.03	90	286.03	90	86.8	90	
	1000	0.20	48.80	56	0.02	104.67	111.05	90	0.53	65.07	33.6	<b>0.01</b>	90.07	85.7	<b>0.01</b>	90.07	85.7	571.88	90	571.88	90	85.7	90	
ta	60	-	-	-	0.22	0.12	0.15	90	0.09	0.78	88.4	<b>0.00</b>	1.85	90.0	<b>0.00</b>	1.85	90.0	21.87	90	21.87	90	90.0	90	
	120	-	-	-	0.71	0.62	1.39	88	1.18	3.41	49.4	<b>0.05</b>	23.34	88.1	<b>0.05</b>	23.34	88.1	41.39	90	41.39	90	88.1	90	
	249	-	-	-	0.66	3.84	40	<b>0.29</b>	<b>0.15</b>	258.18	48	0.57	32.87	44.1	0.40	128.30	44.5	167.61	70	167.61	70	44.5	70	
	501	-	-	-	0.34	36.39	23	<b>0.15</b>	<b>0.00</b>	0.69	90	0.52	1.37	66.8	<b>0.00</b>	3.26	90.0	49.27	90	49.27	90	90.0	90	
ua	120	-	-	-	0.18	0.55	82	<b>0.00</b>	<b>0.00</b>	0.69	90	0.52	1.37	66.8	<b>0.00</b>	3.26	90.0	49.27	90	49.27	90	90.0	90	
	250	-	-	-	0.13	3.53	77	<b>0.02</b>	<b>0.02</b>	5.91	87	0.55	6.83	58.9	<b>0.02</b>	28.70	87.1	100.64	89	100.64	89	87.1	89	
	500	-	-	-	0.08	26.69	68	<b>0.01</b>	<b>0.01</b>	60.04	80	0.38	26.62	60.8	0.08	154.09	69.8	200.77	82	200.77	82	69.8	82	
	1000	-	-	-	0.05	195.11	67	<b>0.01</b>	<b>0.01</b>	491.70	79	0.27	55.74	67.6	0.15	399.37	67.8	400.03	82	400.03	82	67.8	82	
d	120	-	-	-	-	-	-	-	0.04	1.01	88.1	<b>0.00</b>	3.17	90.0	<b>0.00</b>	3.17	90.0	61.82	90	61.82	90	90.0	90	
	250	-	-	-	-	-	-	-	0.08	2.77	80.6	<b>0.00</b>	11.57	89.9	<b>0.00</b>	11.57	89.9	127.93	90	127.93	90	89.9	90	
	500	-	-	-	-	-	-	-	0.10	8.96	69.2	<b>0.00</b>	23.35	89.5	<b>0.00</b>	23.35	89.5	252.79	90	252.79	90	89.5	90	
	120	-	-	-	2.04	0.94	37	<b>0.12</b>	<b>0.12</b>	5.73	63	0.98	1.18	55.0	0.80	12.91	57.7	23.63	67	23.63	67	57.7	67	
da	250	-	-	-	1.13	7.27	28	<b>0.04</b>	<b>0.04</b>	56.71	49	1.32	4.88	38.0	1.14	25.50	40.5	44.66	50	44.66	50	40.5	50	
	500	-	-	-	0.63	70.14	28	<b>0.10</b>	<b>0.10</b>	515.23	47	1.85	18.12	37.0	1.83	53.31	37.1	84.12	49	84.12	49	37.1	49	
	Avg (0)+(u)	0.31	26.37	542	0.18	20.57	600	0.00	23.71	719	0.34	13.60	495.2	0.02	24.56	702.9	197.88	720	197.88	720	702.9	720	702.9	720
	Avg All-(d)	-	-	-	0.40	26.83	1199	0.04	84.63	1507	0.60	14.39	1103.5	0.27	58.01	1426.8	147.41	1567	147.41	1567	147.41	1567	1426.8	1567



**Table 2** Comparison of the current state-of-the-art BPPC heuristics for each instance set and density  $\rho$

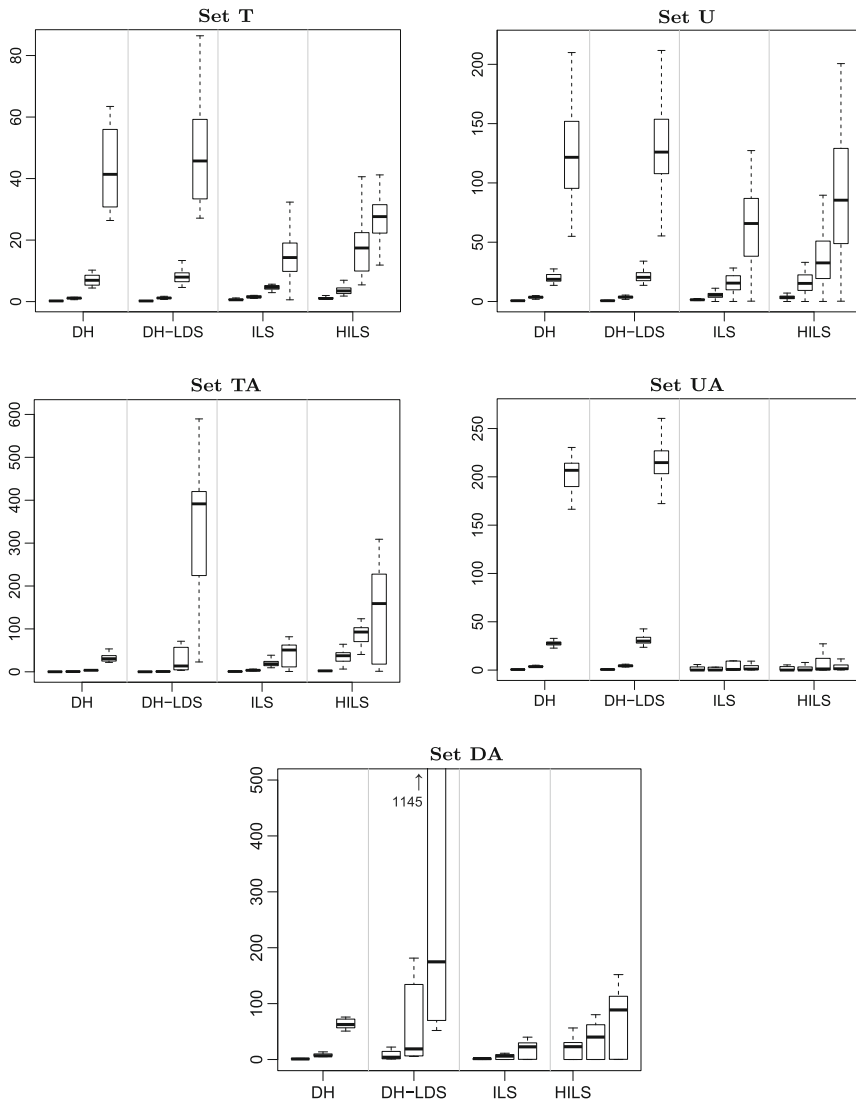
Instance	Class	$\rho$	PH			DH			DH-LDS			ILS-Simple			HILS-Complete			BKS	
			Gap	T(s)	Opt	Gap	T(s)	Opt	Gap	T(s)	Opt	Gap	T(s)	Opt	Gap	T(s)	Opt	Bins	Opt
t		0	<b>0.00</b>	0.00	40	-	-	-	-	-	-	1.44	0.97	25.3	<b>0.00</b>	0.93	40.0	77.50	40
		10	2.08	62.09	2	0.54	5.57	26	<b>0.00</b>	11.35	40	1.26	1.12	26.1	<b>0.00</b>	1.05	40.0	77.55	40
		20	1.20	61.15	9	0.69	6.64	19	<b>0.00</b>	20.87	40	0.72	1.67	29.4	<b>0.00</b>	1.77	40.0	77.73	40
		30	0.71	53.30	18	0.55	9.30	24	<b>0.06</b>	22.98	39	0.45	7.88	27.2	0.21	19.99	32.8	78.33	40
		40	0.14	11.37	36	0.27	16.88	34	<b>0.00</b>	18.36	40	0.23	4.09	30.0	<b>0.00</b>	8.78	40.0	94.10	40
		50	<b>0.00</b>	5.90	40	0.23	16.70	35	<b>0.00</b>	17.78	40	0.29	5.00	27.9	<b>0.00</b>	11.35	40.0	117.93	40
		60	<b>0.00</b>	10.18	40	0.15	12.27	34	<b>0.00</b>	13.49	40	0.28	5.87	28.3	<b>0.00</b>	13.10	40.0	141.53	40
		70	<b>0.00</b>	7.46	40	0.04	9.35	37	<b>0.00</b>	9.81	40	0.27	6.89	23.6	<b>0.00</b>	15.85	40.0	164.33	40
		80	<b>0.00</b>	4.72	40	0.02	7.28	37	<b>0.00</b>	8.15	40	0.15	7.36	29.3	<b>0.00</b>	16.94	40.0	187.78	40
u		90	<b>0.00</b>	1.92	40	0.01	4.92	39	<b>0.00</b>	5.24	40	0.08	8.17	32.0	<b>0.00</b>	18.22	40.0	210.58	40
		0	0.31	35.19	18	-	-	-	-	-	-	<b>0.00</b>	0.26	40.0	<b>0.00</b>	0.56	40.0	188.53	40
		10	0.11	26.56	31	0.11	25.78	37	<b>0.00</b>	28.00	40	<b>0.00</b>	0.28	40.0	<b>0.00</b>	0.60	40.0	188.53	40
		20	0.10	26.77	32	0.02	28.86	38	<b>0.00</b>	30.10	40	<b>0.00</b>	0.27	40.0	<b>0.00</b>	0.60	40.0	188.53	40
		30	0.27	35.28	25	0.05	35.04	39	<b>0.00</b>	35.45	40	0.12	1.32	36.9	0.05	2.65	38.7	188.53	40
		40	0.79	42.60	14	0.18	51.53	33	<b>0.00</b>	55.03	40	0.39	20.12	24.9	0.09	58.14	34.1	194.03	40
		50	0.10	21.68	34	0.18	44.47	31	<b>0.00</b>	48.32	40	0.62	22.43	18.2	<b>0.00</b>	41.99	39.1	236.80	40
		60	0.06	28.69	34	0.09	32.07	33	<b>0.00</b>	33.16	40	0.36	27.02	21.6	<b>0.00</b>	44.72	39.9	284.80	40
		70	0.06	29.84	32	0.05	25.25	33	<b>0.00</b>	27.44	40	0.45	35.54	18.0	0.01	62.13	38.8	330.43	40
	80	0.03	27.66	36	0.03	22.15	36	<b>0.00</b>	22.63	40	0.35	42.03	18.0	<b>0.00</b>	68.64	39.5	376.73	40	
	90	0.00	17.55	39	0.02	16.18	35	<b>0.00</b>	18.62	40	0.16	47.72	23.8	<b>0.00</b>	55.48	40.0	423.68	40	

Table 2 continued

Instance	$\rho$	PH			DH			DH-LDS			ILS-Simple			HILS-Complete			BKS	
		Gap	T(s)	Opt	Gap	T(s)	Opt	Gap	T(s)	Opt	Gap	T(s)	Opt	Gap	T(s)	Opt	Bins	Opt
ta	10	-	-	-	0.65	16.71	22	<b>0.00</b>	48.76	40	0.74	1.42	29.2	<b>0.00</b>	1.52	40.0	77.73	40
	20	-	-	-	0.75	15.38	19	<b>0.00</b>	40.73	40	0.57	1.84	30.9	<b>0.00</b>	2.46	40.0	77.78	40
	30	-	-	-	0.81	12.06	15	<b>0.11</b>	68.77	36	0.76	6.73	23.4	0.33	16.99	30.9	77.80	40
	40	-	-	-	0.42	10.55	21	<b>0.27</b>	107.29	28	0.28	13.58	30.1	0.28	34.60	30.5	78.03	40
	50	-	-	-	0.42	8.74	21	<b>0.42</b>	114.98	21	<b>0.42</b>	25.38	21.0	<b>0.42</b>	74.85	21.0	78.03	40
	60	-	-	-	0.06	8.11	32	<b>0.06</b>	76.20	32	<b>0.06</b>	20.04	32.0	<b>0.06</b>	93.88	32.0	78.40	36
ua	70	-	-	-	0.18	6.79	30	<b>0.00</b>	82.11	32	0.10	16.37	30.9	<b>0.00</b>	107.48	32.0	78.50	32
	80	-	-	-	0.42	7.70	34	<b>0.12</b>	8.94	38	1.66	16.72	10.5	0.44	82.58	20.3	79.00	40
	90	-	-	-	0.66	6.15	18	<b>0.13</b>	87.27	26	1.66	10.27	16.2	0.59	100.13	27.2	82.75	30
	10	-	-	-	0.01	66.60	39	<b>0.00</b>	71.09	40	0.02	1.14	39.0	0.01	2.46	39.3	187.58	40
	20	-	-	-	0.01	64.29	39	<b>0.00</b>	70.83	40	0.01	0.48	39.6	0.01	0.81	39.7	188.35	40
	30	-	-	-	0.05	63.81	38	<b>0.00</b>	65.73	40	<b>0.00</b>	0.40	40.0	<b>0.00</b>	0.71	40.0	187.90	40
	40	-	-	-	0.07	61.56	37	<b>0.00</b>	66.34	40	0.06	1.79	38.7	<b>0.00</b>	3.69	39.7	186.78	40
	50	-	-	-	0.08	60.53	37	<b>0.00</b>	63.45	40	0.05	1.72	39.0	<b>0.00</b>	4.07	40.0	186.93	40
	60	-	-	-	0.18	56.36	33	<b>0.00</b>	77.56	40	0.18	8.54	34.1	0.02	13.52	38.6	187.03	40
	70	-	-	-	0.20	52.10	28	<b>0.03</b>	101.67	38	0.54	39.83	18.0	0.10	75.84	30.7	187.48	39
	80	-	-	-	0.18	45.51	25	<b>0.00</b>	163.48	34	1.08	86.13	2.0	0.18	306.37	21.8	187.75	34
	90	-	-	-	0.25	37.51	18	<b>0.07</b>	576.12	24	1.91	63.73	3.7	0.26	909.73	24.9	189.33	30


Table 2 continued

Instance	$\rho$	PH			DH			DH-LDS			ILS-Simple			HILS-Complete			BKS	
		Gap	T(s)	Opt	Gap	T(s)	Opt	Gap	T(s)	Opt	Gap	T(s)	Opt	Gap	T(s)	Opt	Bins	Opt
d	10	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	43.90	30
	20	-	-	-	-	-	-	-	-	0.06	2.52	28.7	0.00	0.00	5.22	30.0	58.67	30
	30	-	-	-	-	-	-	-	-	0.13	3.13	26.1	0.01	0.01	9.31	29.5	87.67	30
	40	-	-	-	-	-	-	-	-	0.17	3.57	25.8	0.00	0.00	10.70	30.0	114.30	30
	50	-	-	-	-	-	-	-	-	0.02	4.05	28.6	0.00	0.00	13.59	30.0	147.63	30
	60	-	-	-	-	-	-	-	-	0.11	4.82	24.0	0.00	0.00	16.14	30.0	174.07	30
	70	-	-	-	-	-	-	-	-	0.11	5.81	22.6	0.00	0.00	17.90	29.9	204.40	30
	80	-	-	-	-	-	-	-	-	0.03	6.50	26.8	0.00	0.00	19.58	30.0	235.07	30
	90	-	-	-	-	-	-	-	-	0.04	7.83	25.3	0.00	0.00	21.82	30.0	261.93	30
da	10	-	-	-	1.15	31.13	20	0.04	88.01	29	0.00	30.0	0.00	0.00	0.00	30.0	43.90	30
	20	-	-	-	1.25	29.59	16	0.09	80.40	29	0.00	30.0	0.00	0.00	0.00	30.0	43.80	30
	30	-	-	-	1.98	26.21	16	0.00	34.93	30	0.38	28.0	0.38	0.38	3.29	28.0	44.17	30
	40	-	-	-	1.56	22.81	16	0.00	75.94	28	0.69	22.4	0.55	0.55	18.54	23.6	44.03	28
	50	-	-	-	0.91	20.28	18	0.00	77.07	24	2.40	10.05	1.6	1.97	67.05	4.6	44.43	24
	60	-	-	-	1.01	24.36	0	0.04	365.78	2	2.65	12.14	0.0	2.30	74.63	0.0	45.50	3
	70	-	-	-	1.18	43.66	1	0.04	517.06	3	3.27	15.25	2.0	3.12	47.54	2.1	49.77	3
	80	-	-	-	1.52	25.24	1	0.21	331.36	7	2.52	16.01	6.0	2.43	35.15	7.0	60.30	8
	90	-	-	-	0.84	11.79	5	0.34	162.44	7	0.57	14.78	10.0	0.55	28.96	10.0	81.33	10



**Fig. 5** CPU time, in seconds, of HILS-Complete and DH-LDS. One boxplot for each group  $\times$  instance size in terms of number of items

It is important to note that DH and DH-DLS exploit column generation and perform a restricted search in the branch-and-bound tree. Despite their good solution quality, these methods suffer from the same limitations as most mathematical-programming algorithms: their computational effort becomes quickly impracticable as  $n$  grows, as noted in Sadykov and Vanderbeck (2013): “a disadvantage of our primal heuristic is that the running time increases rapidly with the number of items”. To compare the scalability of the methods, Fig. 5 reports the CPU time of DH, DH-DLS, ILS-Simple




**Fig. 6** HILS-Complete GAP by each group of instance

and HILS-Complete for the instances with densities  $\rho \in [30, 70]$ . The CPU time of DH-DLS is initially lower than that of HILS for very small instances, but it tends to increase by a factor 10 every time the instance size doubles. This effect is especially visible on the instances of set (da), which include a larger number of items per bin. On these instances, the column generation quickly becomes a bottleneck. We fitted the CPU time as a power law  $f(n) = \alpha \times n^\beta$  of the number of items  $n$  for each instance class (least-squares regression of an affine function on the log-log graph). The observed CPU time of HILS-Complete grows as  $\mathcal{O}(n^{1.49})$ . In contrast, the time of DH-DLS grows at a rate which is faster than cubic for some instance sets, with power laws of the form  $n^\beta$  where  $2.45 \leq \beta \leq 3.62$ .

Finally, another disadvantage of column generation heuristics relate to the absence of *anytime behavior*: for large instances with one thousand of items, the root node resolution can take several minutes or even hours. Interrupting the search before this time does not allow to retrieve a primal or dual solution. In contrast, a metaheuristic such as HILS can return a feasible primal solution at any time, and be scaled up to solve problem instances of larger size or with a larger number of items per bin. Therefore, even if HILS-Complete does not directly outperform DH-LDS in terms of solution quality, it provides a meaningful alternative for larger or dynamic problems, which are prevalent in a number of emerging applications related to the field of cloud computing and social networks analysis. Moreover, having both efficient metaheuristics and mathematical-programming techniques can open the way to future method hybridizations (*matheuristics*) to harness the strengths of both approaches.

In a second analysis, we evaluate the impact of some instance characteristics on the performance of HILS-Complete. Figure 6 displays the boxplots of the percentage gap for each combination of class  $\times$  number of items. For the data sets (t), (u), and (d) with interval conflict graph, the optimal solution was attained by HILS-Complete on nearly all instances. As a consequence, the boxplots are represented as a single line at a gap of 0.00% and any other gap value appears as an outlier. The other data sets, especially (ta) and (da) with arbitrary conflict graph, are more difficult, such that a gap of 0.00% does not represent the average behavior of the method anymore. The difficulty of the




**Fig. 7** HILS-Complete GAP by density of the conflict graph

instances tends to grow with their size, although this effect is counterbalanced by the way the “gap” metric is defined: on small instances with few bins in the optimal solution, a *mistake* of one bin directly translates into a large gap (e.g., 5% if the BKS includes 20 bins) leading to some outliers for small problems.

Instance class and size are not the only factors which impact the performance of the methods. The density of the conflict graph is also important. This effect is illustrated in Fig. 7, which displays the boxplots of the percentage gap as a function of density. Generally, the most difficult instances are neither those with low density (equivalent to a bin packing problem), nor those with high density (equivalent to a tightly-constrained coloring problem), but those which combine the effects of both packing and conflict constraints at a density level between 30 and 80%. HILS-Complete appears to perform very well for problem instances with low density (up to 20%). On these instance, the BKS has been systematically reached and even sometimes improved, as highlighted by the two outliers located below 0.00%.

We also evaluated the impact of the large neighborhoods, comparing the results of ILS-Simple and HILS-Complete. We observe a significant difference of solution quality, confirmed by a pairwise Wilcoxon test ( $p$  value  $\leq 10^{-64}$ ). The large neighborhoods have a positive impact on solution quality, but consume some additional CPU time (49.85 s in average per instance with the large neighborhoods, and 12.53 s without). Note that we also conducted side experiments to investigate whether an increase in the number of iterations of ILS-Simple could help to profit from this additional CPU time to reach better solutions, but the quality improvements were only minor when doubling the number of iterations (see Sect. 4.2).

Finally, we analyzed the time consumption of each component of the method. The results of this analysis are reported in Fig. 8, which displays the percentage CPU time of each main component of the approach. This metric has been aggregated per instance class  $\times$  number of items. The printed article contains the figure in black and white, and we refer to the on-line article for a color figure. The search effort appears to be well distributed among the search components, i.e., no specific neighborhood consumes the majority of the time. Overall, the assignment neighborhood, local search



**Fig. 8** Time consumption of each component of HILS-Complete

and zero-cost moves are the largest time consumers. In contrast, the ejection chains and grenade neighborhoods use less than 6 and 4% of the CPU time, respectively. The computational burden of the 0-cost moves mostly relates to the fact that it constitutes the first loop of the local search, and thus involves many moves evaluations which are later known as non-improving, and thus pruned in subsequent local search iterations. In line with the previous discussions, we also observe that the neighborhood based on set covering requires a larger amount of CPU time for difficult instances with an arbitrary conflict graph, such as (ta) and (da).

Overall, from these experiments, HILS-Complete is a promising alternative method for the BPPC, producing high quality solutions with an overall gap of 0.22% from the BKS, and improving upon the previous best metaheuristic for the problem in a more controlled CPU time than existing mathematical programming-based methods. Having two alternative methods with their own strengths (near optimality for DH-LDS, and good scalability for HILS) will be a good asset for future works, which can explore new hybridizations. Since the neighborhoods are central in the design of HILS, we now go a step further to identify the most promising search components and measure the impact of its key parameters. This is the subject of a detailed analysis in the next section.

#### 4.2 Sensitivity analysis: search parameters and neighborhoods

This section analyses the sensibility of the method to a change of its parameters, and evaluates the contribution of its key components and neighborhoods. Starting from the standard HILS-Complete configuration, we varied one parameter/factor at a time (OFAT approach) and tested the resulting configurations. These experiments are summarized in Tables 3, 4 and 5. Each line presents the performance of an alternative configuration of the method (indexed from A to R), in terms of average gap and time, on the six classes of instances. To maintain a reasonable computational effort, a single run was performed for each configuration on the 2060 instances.

**Table 3** Impact of some variations of the shaking and termination parameters

	$N_{\text{SHAK}}$	$N_{\text{LS}}$	$S_{\text{SHAK}}$	(t)		(u)		(ta)		(ua)		(d)		(da)	
				Gap	Time	Gap	Time	Gap	Time	Gap	Time	Gap	Time	Gap	Time
Standard	50	100	3	<b>0.02</b>	9.97	0.01	34.44	0.23	52.22	0.05	190.40	<b>0.00</b>	11.78	1.16	31.45
A. More frequent	100	50	3	<b>0.02</b>	13.76	0.01	42.16	0.22	80.40	0.04	250.58	<b>0.00</b>	13.81	1.16	57.97
B. Less frequent	25	200	3	0.04	10.96	0.02	32.58	0.26	46.08	0.10	73.39	<b>0.00</b>	13.05	1.40	22.70
C. Longer run	100	100	3	<b>0.02</b>	21.03	<b>0.00</b>	67.83	0.21	97.67	<b>0.03</b>	382.53	<b>0.00</b>	25.35	<b>1.09</b>	59.95
D. Shorter run	25	100	3	0.04	5.65	0.02	20.18	0.28	34.69	0.10	50.05	0.01	6.45	1.47	15.60
E. Stronger shaking	50	100	5	0.03	10.93	0.01	37.71	0.24	56.82	0.09	111.00	<b>0.00</b>	12.87	1.17	31.18
F. Weaker shaking 1	50	100	1	<b>0.02</b>	10.56	0.01	31.57	<b>0.18</b>	76.25	0.06	166.64	<b>0.00</b>	12.47	1.42	25.78
G. Weaker shaking 2	50	100	2	0.03	11.03	0.01	32.81	0.24	62.01	0.07	132.53	<b>0.00</b>	12.73	1.29	30.31



**Table 4** Sensitivity analysis when deactivating some neighborhoods

	(t)		(u)		(ta)		(ua)		(d)		(da)	
	Gap	Time	Gap	Time	Gap	Time	Gap	Time	Gap	Time	Gap	Time
Standard	<b>0.02</b>	9.97	<b>0.01</b>	34.44	<b>0.23</b>	52.22	<b>0.05</b>	190.40	<b>0.00</b>	11.78	<b>1.16</b>	31.45
H. No assign	0.03	5.21	<b>0.01</b>	24.38	0.27	43.48	0.10	61.89	<b>0.00</b>	4.73	1.29	28.37
I. No ejection chains	<b>0.02</b>	9.99	<b>0.01</b>	30.92	0.24	56.67	0.07	108.31	<b>0.00</b>	11.58	1.23	28.87
J. No grenade	<b>0.02</b>	9.84	0.02	34.99	0.25	46.49	0.07	70.97	0.01	11.58	1.21	29.51
K. No set covering	0.37	9.28	0.03	33.16	0.54	33.76	0.30	46.34	<b>0.00</b>	11.88	1.40	13.75
L. No large neighborhoods	0.49	4.50	0.24	18.66	0.72	11.35	0.43	18.98	0.07	3.97	1.36	7.39
M. No large neighborhoods ++	0.39	50.21	0.15	65.62	0.65	27.61	0.39	47.52	0.05	15.56	1.22	16.33
N. No 0-cost moves	0.10	9.52	0.02	32.63	0.31	57.97	0.07	172.05	<b>0.00</b>	10.15	3.16	17.06

**Table 5** Sensitivity analysis on the parameters of the set covering neighborhood

	Adapt	$S_{pool}$	$T_{LMIT}$	(t)		(u)		(ta)		(ua)		(d)		(da)	
				Gap	Time	Gap	Time	Gap	Time	Gap	Time	Gap	Time	Gap	Time
Standard	Yes	1500	20	<b>0.02</b>	9.97	<b>0.01</b>	34.44	0.23	52.22	<b>0.05</b>	190.40	<b>0.00</b>	11.78	<b>1.16</b>	31.45
O. Static pool -	No	1000	10	0.05	9.80	0.02	34.14	0.33	41.57	0.16	44.97	<b>0.00</b>	11.84	1.30	16.58
P. Static	No	1500	20	0.04	9.85	0.02	35.18	0.26	54.75	0.14	48.92	<b>0.00</b>	11.80	1.20	24.84
Q. Static pool +	No	2000	60	0.03	11.42	0.02	34.29	<b>0.20</b>	95.66	0.11	71.90	<b>0.00</b>	11.93	1.28	58.21
R. Set partitioning	Yes	1500	20	<b>0.02</b>	10.19	<b>0.01</b>	33.85	0.22	52.82	0.07	122.83	<b>0.00</b>	12.13	1.20	27.71

*Impact of the general search parameters.* Our first investigation concerns the shaking operator and the termination criterion, as these parameters are recurrent in all metaheuristics based on iterated local search. In the first two alternative configurations, we modified the periodicity of the shaking operator, i.e., the number  $N_{LS}$  of zero-cost moves and local search phases before shaking. This parameter is halved in configuration A, and doubled in configuration B. Since the product  $N_{SHAK} \times N_{LS}$  governs the total number of local search descents before termination, we kept it constant by updating  $N_{SHAK}$  accordingly. The configurations C and D aim to investigate the impact of longer and shorter runs, respectively, by doubling or halving the value  $N_{SHAK}$  without impacting  $N_{LS}$ . Finally, the configurations E, F and G investigate alternative strengths for the shaking operator, which may involve to perform stronger shaking for a better exploration of the search space ( $S_{SHAK} = 5$ ), or weaker shaking for further intensification ( $S_{SHAK} = 1$  or 2).

The changes of parameter setting for the configurations A, B, E, F and G are generally detrimental for the performance of the method. It appears that strengthening the shaking, or calling upon this operator more frequently has only a small impact on performance, but the inverse process would have a larger negative impact. The configuration D allows a longer run, and thus naturally leads to better solutions. This improvement of solution quality, however, remains minor in comparison to the additional CPU time.

*Contribution of the large neighborhoods.* The second experiment aims to evaluate the respective role of each large neighborhood as well as the 0-cost moves, which are critical for search diversification. We thus deactivated, in the configurations H to K, each one of the four large neighborhoods. In configuration L, we deactivated all four neighborhoods. Finally, no 0-cost moves are applied in configuration N.

All search components appear to contribute positively to the final solution quality. By order of magnitude, the largest loss occurs when deactivating the 0-cost moves, as they contribute very significantly to diversify the search. In second position, deactivating all large neighborhoods (as in *ILS-Simple*) also leads to a large loss of solution quality, but also to a gain of computational time. We tested longer runs (configuration M, via a fourfold increase of  $N_{SHAK}$ ) without the large neighborhoods, and this was not sufficient to regain the solution quality of the complete method.

The large neighborhood based on set covering has the largest impact on the search, especially for the instance classes (t) and (ua). As noted in Fig. 8, on these particular instances, the set covering problem is conveniently solved in small CPU time. This allows to use a larger pool of columns, thus increasing the chances of solution improvement. Finally, the last three neighborhoods, ejection chains, assignment, and grenade have a more moderate impact on solution quality. Their contributions are only observed on the difficult instances: (ta), (ua) and (da). Nevertheless, note that the removal of these three neighborhoods together is more largely detrimental (compare configurations K and L). The three neighborhoods provide different forms of solutions improvement and diversification, which are not critical when taken one by one, but much more significant when considered as a whole.

*Impact of the set-covering parameters.* The last parameters of the method concern the set covering neighborhood. Our preliminary calibration led to a time limit  $T_{LIMIT} = 20$  seconds for the solver, as well as an adaptive mechanism for the size

of the pool, starting from a pool size of  $S_{\text{POOL}} = 1500$  columns. We thus tested the impact of this adaptive mechanism in the configurations O to Q, by evaluating three static pool size values of 1000, 1500 and 2000 columns, along with a time of 10, 20 and 60 seconds, respectively, for the resolution of the set covering problem. Finally, the last configuration R considers a possible replacement of the set covering by a set partitioning formulation, using an equality in Eq. (11).

In these experiments, we observe that the size of the pool is a very sensible parameter. A small value leads to simple set covering problems which are not likely to result in solution improvements, even when solved to optimality, while a larger value opens the way to more improvement opportunities, at the risk of not solving the mathematical model in the allowed time. For the instances of class (da), an intermediate pool size of 1500 appears to be desirable, while for other instances, e.g., (ta) and (ua), a larger pool size with 2000 columns leads to solutions of better quality at the cost of additional computational time. Clearly, the best parameter setting is dependent on the instance characteristics. This led us to use the adaptive mechanism, which performs better overall than each of the static configurations. Finally, using either the set covering or the set partitioning model did not lead to a significant difference of solution quality.

## 5 Conclusions

In this paper, we introduced an ILS based on several classes of local and large (exponential-sized) neighborhoods for the BPPC. We intentionally relied on simple metaheuristic concepts so as to draw the complete focus on the large neighborhoods and accurately measure their contribution to the search. We proposed  $\mathcal{O}(1)$  move evaluation procedures for the local search, polynomial variants of ejection chains and assignment neighborhoods, as well as an adaptive search mechanism for the set covering-based neighborhood. We also introduced a controlled use of 0-cost moves to further diversify the search.

The resulting ILS produces solution of good quality on the classic BPPC benchmark instances, with an average gap to the BKS of 0.22%. This is a better solution quality than the previous best metaheuristic, the genetic algorithm of Fernandes Muritiba et al. (2010), and a slightly lower quality than the diving heuristic with limited discrepancy search (DH-LDS) of Sadykov and Vanderbeck (2013). Still, the current instances are an ideal testing ground for column-generation-based methods, e.g. DH-LDH, as they contain very few items per bin (2.87 in average). On the other hand, the proposed metaheuristic exploits different search concepts, and scales very well with an increase of problem size, with a measured CPU time in  $\mathcal{O}(n^{1.49})$ . In comparison, the CPU time of DH-LDH tends to increase by a factor ten for each twofold increase of problem size. Finally, we conducted extensive experiments to measure the contribution of each neighborhood as well as the impact of the key search parameters. As underlined by these experiments, the 0-cost moves are critical for the search performance, as well as the set covering-based neighborhood. The other large neighborhoods, i.e., ejection chains, assignment, and grenade, have a smaller individual impact but contribute significantly to the search when used together.

Many research perspectives are open on this class of problems. First, future hybridizations should allow to combine the strengths of both neighborhood-centered search and column-generation based-heuristics, with the aim of keeping both near-optimality and tractability. Future research should thus be done along the path of matheuristics, aiming to harness the strength of mathematical programming and neighborhood search. Second, populations of solutions and more advanced memory structures can be a strong asset to help diversifying the search and favoring the discovery of better solutions. Finally, as the current benchmark instances are now solved to near-optimality, new and larger test sets, in terms of overall number of items, and number of items per bin, should also be investigated.

**Acknowledgements** This research was partially supported by the Brazilian agencies CAPES, CNPq and FAPERJ.

## References

- Ahuja, R.K., Ergun, O., Orlin, J.B., Punnen, A.P.: A survey of very large-scale neighborhood search techniques. *Discrete Appl. Math.* **123**(1–3), 75–102 (2002)
- Avanthay, C., Hertz, A., Zufferey, N.: A variable neighborhood search for graph coloring. *Eur. J. Oper. Res.* **151**(2), 379–388 (2003)
- Blöchliger, I., Zufferey, N.: A graph coloring heuristic using partial solutions and a reactive tabu scheme. *Comput. Oper. Res.* **35**(3), 960–975 (2008)
- Blum, C., Roli, A.: Metaheuristics in combinatorial optimization: overview and conceptual comparison. *ACM Comput. Surv.* **35**(3), 268–308 (2003)
- Deineko, V.G., Woeginger, G.J.: A study of exponential neighborhoods for the travelling salesman problem and for the quadratic assignment problem. *Math. Program.* **87**(3), 519–542 (2000)
- Delorme, M., Iori, M., Martello, S.: Bin packing and cutting stock problems: mathematical models and exact algorithms. *Eur. J. Oper. Res.* **255**(1), 1–20 (2016)
- Dowland, K., Thompson, J.: An improved ant colony optimisation heuristic for graph colouring. *Discrete Appl. Math.* **156**(3), 313–324 (2008)
- Elhedhli, S., Li, L., Gzara, M., Naoum-Sawaya, J.: A branch-and-price algorithm for the bin packing problem with conflicts. *Inf. J. Comput.* **23**(3), 404–415 (2011)
- Epstein, L., Levin, A.: On bin packing with conflicts. *SIAM J. Optim.* **19**(3), 1270–1298 (2008)
- Epstein, L., Levin, A., Stee, R.: Two-dimensional packing with conflicts. *Acta Inf.* **45**(3), 155–175 (2008)
- Falkenauer, E.: A hybrid grouping genetic algorithm for bin packing. *J. Heuristics* **2**(1), 5–30 (1996)
- Fernandes Muritiba, A., Iori, M., Malaguti, E., Toth, P.: Algorithms for the bin packing problem with conflicts. *Inf. J. Comput.* **22**(3), 401–415 (2010)
- Fleszar, K., Charalambous, C.: Average-weight-controlled bin-oriented heuristics for the one-dimensional bin-packing problem. *Eur. J. Oper. Res.* **210**(2), 176–184 (2011)
- Galiner, P., Hao, J.: Hybrid evolutionary algorithms for graph coloring. *J. Comb. Optim.* **3**(4), 379–397 (1999)
- Gendreau, M., Laporte, G., Semet, F.: Heuristics and lower bounds for the bin packing problem with conflicts. *Comput. Oper. Res.* **31**(3), 347–358 (2004)
- Glover, F.: Ejection chains, reference structures and alternating path methods for traveling salesman problems. *Discrete Appl. Math.* **65**(1–3), 223–253 (1996)
- Gutin, G.: Exponential neighbourhood local search for the traveling salesman problem. *Comput. Oper. Res.* **26**(4), 313–320 (1999)
- Hamdi-Dhaoui, K., Labadie, N., Yalaoui, A.: The bi-objective two-dimensional loading vehicle routing problem with partial conflicts. *Int. J. Prod. Res.* **52**(19), 5565–5582 (2014)
- Hertz, A., de Werra, D.: Using tabu search techniques for graph coloring. *Computing* **39**(4), 345–351 (1987)
- Jansen, K.: An approximation scheme for bin packing with conflicts. *J. Comb. Optim.* **3**(4), 363–377 (1999)
- Jansen, K., Öhring, S.: Approximation algorithms for time constrained scheduling. *Inf. Comput.* **132**(2), 85–108 (1997)

- Khanafar, A., Clautiaux, F., Talbi, E.: New lower bounds for bin packing problems with conflicts. *Eur. J. Oper. Res.* **206**(2), 281–288 (2010)
- Kuhn, H.W.: The hungarian method for the assignment problem. *Naval research logistics quarterly* **2**(1–2), 83–97 (1955)
- Laporte, G., Desroches, S.: Examination timetabling by computer. *Comput. Oper. Res.* **11**(4), 351–360 (1984)
- Lewis, R.: *A Guide to Graph Colouring: Algorithms and Applications*. Springer, Berlin (2016)
- Lewis, R., Thompson, J., Mumford, C., Gillard, J.: A wide-ranging computational comparison of high-performance graph colouring algorithms. *Comput. Oper. Res.* **39**(9), 1933–1950 (2012)
- Malaguti, E., Toth, P.: A survey on vertex coloring problems. *Int. Trans. Oper. Res.* **17**(1), 1–34 (2010)
- Malaguti, E., Monaci, M., Toth, P.: A metaheuristic approach for the vertex coloring problem. *Inf. J. Comput.* **20**(2), 302–316 (2008)
- Masson, R., Vidal, T., Michallet, J., Penna, P.H.V., Petrucci, V., Subramanian, A., Dubedout, H.: An iterated local search heuristic for multi-capacity bin packing and machine reassignment problems. *Expert Syst. Appl.* **40**(13), 5266–5275 (2013)
- Minh, T.T., Hoai, T.V., Nguyet, T.T.N.: A memetic algorithm for waste collection vehicle routing problem with time windows and conflicts. In: Murgante, B. (ed.) *Computational Science and Its Applications—ICCSA 2013*, pp. 485–499. Springer, Heidelberg (2013)
- Monaci, M., Toth, P.: A set-covering-based heuristic approach for bin-packing problems. *Inf. J. Comput.* **18**(1), 71–85 (2006)
- Morgenstern, C.: Distributed coloration neighborhood search. In: Johnson, D., Trick, M. (eds.) *Discrete Mathematics and Theoretical Computer Science*, pp. 335–358. American Mathematical Society, Providence (1996)
- Muter, I., Birbil, S.I., Sahin, G.: Combination of metaheuristic and exact algorithms for solving set covering-type optimization problems. *Inf. J. Comput.* **22**(4), 603–619 (2010)
- Quiroz-Castellanos, M., Cruz-Reyes, L., Torres-Jimenez, J., Gómez, C., Huacuja, H.J.F., Alvim, A.C.: A grouping genetic algorithm with controlled gene transmission for the bin packing problem. *Comput. Oper. Res.* **55**, 52–64 (2015)
- Sadykov, R., Vanderbeck, F.: Bin packing with conflicts: a generic branch-and-price algorithm. *Inf. J. Comput.* **25**(2), 244–255 (2013)
- Sarvanov, V., Doroshko, N.: The approximate solution of the travelling salesman problem by a local algorithm with scanning neighborhoods of factorial cardinality in cubic time (in Russian). *Software: Algorithms and Programs 31*, pp. 11–13. Mathematical Institute of the Belarusian Academy of Sciences, Minsk (1981)
- Sörensen, K.: Metaheuristics—the metaphor exposed. *Int. Tran. Oper. Res.* **22**(1), 3–18 (2015)
- Subramanian, A., Uchoa, E., Ochi, L.: A hybrid algorithm for a class of vehicle routing problems. *Comput. Oper. Res.* **40**(10), 2519–2531 (2013)
- Thompson, P., Psaraftis, H.: Cyclic transfer algorithms for multi-vehicle routing and scheduling problems. *Oper. Res.* **41**(5), 935–946 (1993)
- Toth, P., Tramontani, A.: An integer linear programming local search for capacitated vehicle routing problems. In: Golden, B., Raghavan, S., Wasil, E. (eds.) *The Vehicle Routing Problem: Latest Advances and New Challenges*, pp. 275–295. Springer, Boston (2008)
- Vidal, T., Crainic, T., Gendreau, M., Prins, C.: Heuristics for multi-attribute vehicle routing problems: a survey and synthesis. *Eur. J. Oper. Res.* **231**(1), 1–21 (2013)