

Bin Packing Problem with Time Lags

Orlando Rivera Letelier, François Clautiaux, Ruslan Sadykov

▶ **To cite this version:**

Orlando Rivera Letelier, François Clautiaux, Ruslan Sadykov. Bin Packing Problem with Time Lags. INFORMS Journal on Computing, 2022, 34 (4), pp.2249-2270. 10.1287/ijoc.2022.1165 . hal-02986895v2

HAL Id: hal-02986895

https://inria.hal.science/hal-02986895v2

Submitted on 14 Dec 2022

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Bin Packing Problem with Time Lags

Orlando Rivera Letelier*[1], François Clautiaux†[2,3], and Ruslan Sadykov‡[2,3]

[1]Doctoral Program in Industrial Engineering and Operations Research, Universidad Adolfo Ibáñez, Av. Diagonal Las Torres 2640, Peñalolén, Santiago, Chile. 7941169
[2]Inria Bordeaux – Sud-ouest, 200 avenue de la Vieille Tour, 33405 Talence, France
[3]IMB, Université de Bordeaux, 351 cours de la Libération, 33405 Talence, France

2 January 2022

Abstract

We introduce and motivate several variants of the bin packing problem where bins are assigned to time slots, and minimum and maximum lags are required between some pairs of items. We suggest two integer programming formulations for the general problem: a compact one, and a stronger formulation with an exponential number of variables and constraints. We propose a branch-cut-and-price approach which exploits the latter formulation. For this purpose, we devise separation algorithms based on a mathematical characterization of feasible assignments for two important special cases of the problem: when the number of possible bins available at each period is infinite, and when this number is limited to one, and time lags are non-negative. Computational experiments are reported for instances inspired from a real-case application of chemical treatment planning in vineyards, as well as for literature instances for special cases of the problem. The experimental results show the efficiency of our branch-cut-and-price approach, as it outperforms the compact formulation on newly proposed instances, and is able to obtain improved lower and upper bounds for literature instances.

1 Introduction

The bin packing problem is one of the most widely studied combinatorial optimization problems and has been known since the 1930s (Kantorovich, 1960). The classical bin packing problem (BPP) consists in assigning a set of items to a minimum possible number of identical capacitated bins. Formally, we are given a set of items $V = \{1, \ldots, n\}$, and an unlimited quantity of identical bins with a positive capacity W. Each item $i \in V$ has a positive weight $w_i \leq W$. The goal is to find a packing of items into the minimum number of bins, such that the total weight of items packed in each bin does not exceed its capacity.

In this paper, we introduce the *bin packing problem with time lags* (BPPTL), which is a generalization of the BPP. In this variant, items represent tasks that consume resources during one time period. Time lags represent two kinds of constraints: one needs to wait a minimum amount of time between two tasks i and j (classical time lags), or one cannot wait more than a certain amount of time after the end of a task i to process a given task j (this can be modeled by a negative time lag between j and i, as explained below). Here the assignment of items to

*orlandoriveraletelier@gmail.com
†francois.clautiaux@inria.fr
‡ruslan.sadykov@inria.fr

bins is performed over a discretized time horizon, and there are generalized precedence relations between items. In the BPPTL, in addition to the set of items, we have a directed valued graph $G = (V, A, l)$ representing precedence constraints between items, as well as a positive integer value L which defines the upper bound on the number of bins that can be used in each time period. If the value L is large enough, i.e., if in any feasible solution the number of bins is always sufficient at each time period, for instance $L \geq n$, we simply write $L = \infty$. Each arc $(i, j) \in A$ has an associated value $l_{i,j} \in \mathbb{Z}$. The BPPTL consists in assigning each item $i \in V$ to bin $\bar{b}_i \in \{1, \ldots, L\}$ and a time period $\bar{p}_i \in \mathbb{Z}_+$, such that bin capacity (1a) and time lag (1b) constraints are satisfied:

$$\sum_{i \in V:\, \bar{b}_i = b,\, \bar{p}_i = p} w_i \leq W \quad \forall b \in \{1, \ldots, L\},\ \forall p \in \mathbb{Z}_+, \tag{1a}$$

$$\bar{p}_i + l_{i,j} \leq \bar{p}_j \quad \forall (i, j) \in A. \tag{1b}$$

The objective of the BPPTL is to find a feasible assignment that minimizes the number of bin-period pairs (b, p) which have at least one item assigned to them, or to determine that such an assignment does not exist.

Values $l_{i,j}$ and $l_{j,i}$ allow us to model minimal and maximal time lags between a pair of items i and j, i.e., minimal and maximal difference between time periods to which items i and j are assigned. If this difference $\bar{p}_j - \bar{p}_i$ lies inside interval $[\ell^-, \ell^+]$, we set $l_{i,j} = \ell^-$ and $l_{j,i} = -\ell^+$. Therefore, values l can be both positive and negative. By abuse of notation, values l are also called *time lags*.

Our motivation for studying the BPPTL stems from applications in which some tasks should be performed repeatedly using resources that are available on a pay-per-use basis. For example, a given task should be performed six times in a given time period. A possible way to deal with recurrent tasks is to determine an a priori fixed time between two consecutive occurrences of the same task. This over-constrained setting limits the sharing of resources, and may lead to expensive solutions. On the other hand, in many applications the decision maker does not impose a fixed frequency of task occurrences. Only a desired frequency is specified, which may be altered to some extent if this leads to a cost reduction. Thus, a more flexible approach is to impose minimum and maximum time lags between two consecutive occurrences of a task. This allows the decision maker to ensure a better usage of the resource. At the same time, solutions in which tasks are not distributed evenly over the time horizon are forbidden.

A specific case of such a problem can be observed when one performs the planning of phytosanitary treatments in a vineyard. The vineyard is divided into sectors, i.e., well-defined portions of the property that correspond to contiguous geographical areas. One is given a set of meta-requests for treatment, i.e., diseases against which the vineyard must be protected using phytosanitary products. For each appropriate sector-disease pair, a periodic treatment has to be performed. Each treatment consists of a sequence of tasks of a certain duration. A phytosanitary treatment protects the vines only for a certain period of time, so the same task has to be repeated over the time horizon. On the other hand, treatments should not be repeated too often due to regulations on the spread of phytosanitary products. Practically speaking, the same chemical product cannot be spread twice in a given period, whose length depends in the toxicity of the product. Therefore, two consecutive occurrences of the same treatment should be neither too far nor too close in time. Treatments are performed by identical rented vehicles with a fixed cost per day of usage. Vehicles are limited in the total duration of work in a day. A vehicle can perform several treatments in a day as long as their total duration does not exceed the maximum total duration. The time needed for a vehicle to go from one sector to another is negligible in comparison to the duration of treatments. Given the strict regulation on chemical product spreading, tasks are never fragmented and cannot be processed over two time periods.

This application can be modelled as the bin packing problem with time lags. Here, capacitated bins represent vehicles each having a maximum total duration. Items represent treatment tasks of different durations. The time lags graph consists of double-linked chains, each chain representing one periodic treatment. Two consecutive items in the chain represent two consecutive occurrences of a phytosanitary treatment. There are two arcs between consecutive items:

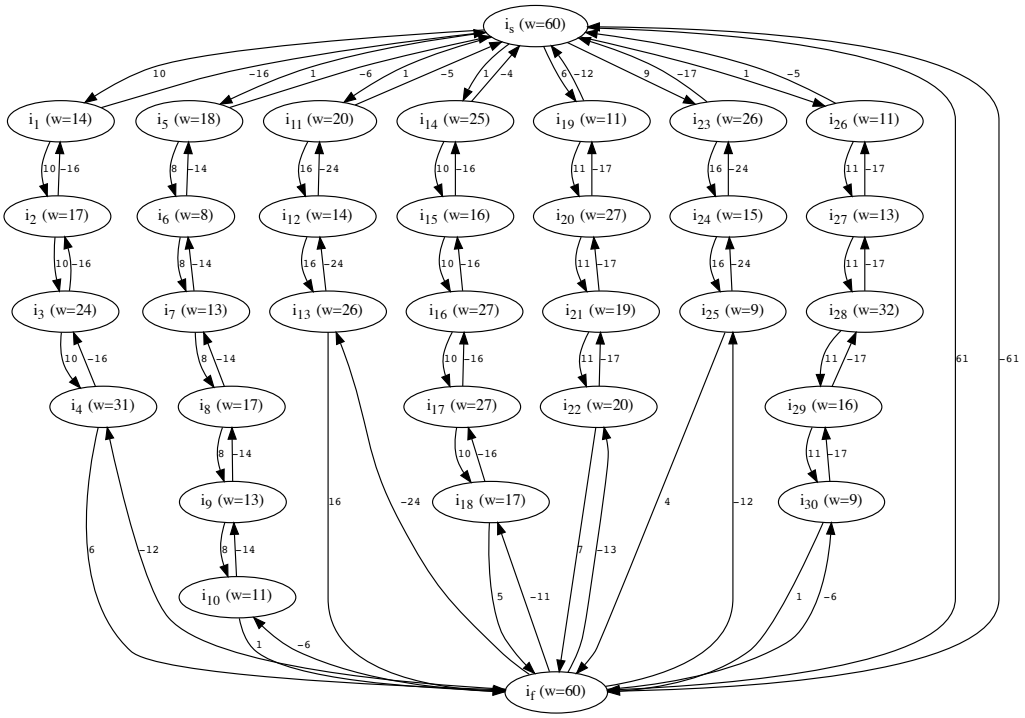


Figure 1: An example of the time lags graph for an instance of the phytosanitary treatments planning application

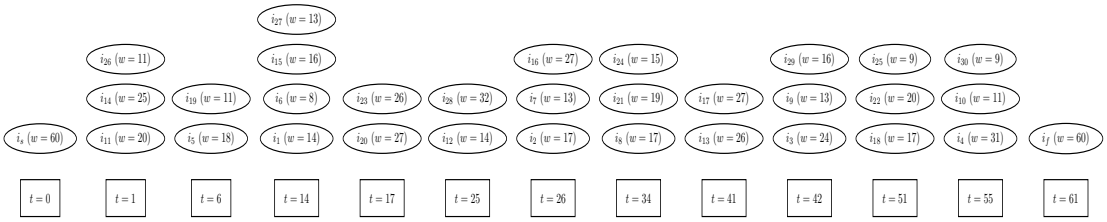


Figure 2: Solution of instance in Figure 1

one with a positive time lag representing the minimum number of days between treatment occurrences, and the other with a negative time lag representing the maximum number of days between occurrences. Two special source and sink nodes together with arcs between them serve to fix the planning time horizon. An example of such graph is depicted in Figure 1. The optimal solution for this example is depicted in Figure 2. It happens that this solution uses at most one bin per time period. In general, however, more than one bin can be used in the same period.

The BPPTL has some similarities with batch scheduling problems, in which jobs are first packed in batches, and then batches are scheduled within a time horizon. There are two main reasons why we attribute our problem to the class of bin packing problems. First, batches in scheduling are typically of different durations, whereas in our case the duration of a bin is always one period. Second, our objective function is typical for bin packing problems (minimizing the number of bins), whereas the total schedule length (makespan) objective is typical for scheduling problems. The difference here is that time periods without items or jobs are penalized in scheduling, but not penalized in the BPPTL.

1.1 Literature review

We now review the literature on the bin packing problems related to the BPPTL. The literature on the classical BPP is vast. Here we only present recent exact approaches to this problem. A comprehensive survey on exact methods for the BPP, including an original comparative computational study, was published by Delorme et al. (2016). After that, other exact specialized algorithms for the BPP were proposed by Delorme and Iori (2020) and by Wei et al. (2020). A generic BCP solver capable of solving exactly the BPP among other problems was proposed by Pessoa et al. (2020).

The BPPTL is closely related to some generalizations of the BPP known in the literature. The bin packing problem with precedence constraints (BPP-P) is a special case of the BPPTL in which all time lags are unitary and $L = 1$. The first exact approach for the BPP-P was developed by Dell'Amico et al. (2012). They proposed a large set of lower bounds, a variable neighborhood search, and a branch-and-bound algorithm. Later, Pereira (2016) suggested a dynamic programming-based heuristic and an exact enumeration procedure which uses several lower bounds and dominance rules to solve the BPP-P.

The simple assembly line balancing problem (SALBP-1) is a special case of the BPPTL in which all time lags are equal to zero and $L = 1$. The survey on simple assembly line balancing problems was presented by Scholl and Becker (2006). Morrison et al. (2014) proposed an exact branch-bound-and-remember algorithm for the SALBP-1. Pereira (2015) analyzed different families of lower bounds for the SALBP-1 and also presented new lower bounds which improved on the best known bounds for open instances of the problem.

Recently a generalization of both the BPP-P and the SALBP-1 was considered by Kramer et al. (2017). They introduced the bin packing problem with generalized precedence constraints (BPP-GP). In the BPP-GP considered by Kramer et al. (2017), time lags can only be non-negative, $L = 1$, and the objective is to minimize the makespan, i.e., the latest time period that has non-empty bins. As shown in Section 4, the number of used bins and the makespan objectives are equivalent for the case when $L = 1$ and every time lag is either one or zero. However, if some time lags are strictly greater than one, then the BPP-GP is not a special case of the BPPTL. In (Kramer et al., 2017), the authors proposed an effective iterated local search algorithm for the BPP-GP with the makespan objective. They also applied preprocessing techniques and lower bounding procedures in order to estimate the quality of the solutions they obtained. Thus, many instances were solved to optimality.

Another related problem is the bin packing problem with conflicts (BPPC), in which a given undirected graph represents conflicts between items. Two adjacent items in the conflict graph cannot be put together in one bin. The objective function is the same as in the standard BPP, i.e., minimization of the number of bins used. The BPPC was introduced by Gendreau et al. (2004). Khanafer et al. (2010, 2012) proposed improved lower bounds for the BPPC based on dual-feasible functions, as well as tree-decomposition-based heuristics. Exact branch-and-price algorithms for the BPPC have been proposed in several papers (Fernandes Muritiba et al., 2010; Elhedhli et al., 2011; Sadykov and Vanderbeck, 2013; Wei et al., 2020).

In branch-and-price algorithms for the BPPC, the pricing problem is the knapsack problem with conflicts (KPC). This problem is interesting to us as it is a special case of the pricing problem considered in Section 3.2.2. The KPC has been solved either by MIP in (Fernandes Muritiba et al., 2010; Elhedhli et al., 2011), by dynamic programming and branch-and-bound algorithms in (Sadykov and Vanderbeck, 2013), or by a labelling algorithm as a resource-constrained shortest path problem in (Wei et al., 2020). Recently, improved combinatorial branch-and-bound algorithms for the KPC were proposed by Bettinelli et al. (2017) and Coniglio et al. (2021).

One more related class of problem is p-batch (or parallel batch) scheduling. In these problems, jobs have processing time and size. A set of jobs can form a batch if their total size does not exceed the batch capacity. When jobs have identical processing times, the p-batch scheduling problem with the makespan objective is equivalent to the standard bin packing problem, as shown by Uzsoy (1994). Uzsoy (1995) considered the p-batch scheduling problem with incompatible

job families, which corresponds to the BPPC. However, only the polynomially solvable variant with identical item sizes was considered. The p-batch scheduling problem with different job processing times and sizes and the makespan objective was considered by Dupont and Dhaenens-Flipo (2002), who proposed an exact branch-and-bound algorithm for it. A branch-and-price algorithm for the same problem with parallel batching machines was suggested by Rafiee Parsa et al. (2010). Among recent papers dealing with batch scheduling problems with precedence constraints we can mention works by Bilyk et al. (2014) and by Emde et al. (2020). We are not aware of works on batch scheduling with generalized time lags. However, in the case of identical job processing times, release dates and deadlines of jobs can be modelled with positive and negative time lags. We would like to conclude this part with a reminder that the main difference between our problem and batch scheduling is the objective function employed. Minimizing the number of bins corresponds to minimizing the number of batches in scheduling. In the presence of precedence constraints, the latter objective is different from ones traditionally used in scheduling like the makespan or tardiness-related criteria.

1.2 Contribution and outline

In this paper, as well as introducing and motivating the BPPTL, we present integer programming (IP) formulations for the problem. The first formulation is a compact one which involves one binary variable for every triple item-bin-time period. The second IP formulation with an exponential number of variables gives a relaxation to the BPPTL, similar to the one used in (Pereira, 2016) for the BPP-P. An exponential set of constraints is then introduced allowing us to turn the second formulation into an exact one for two important special cases of the BPPTL. The first one is when $L = \infty$. The second one is when $L = 1$, and time lags are non-negative. We present two ways of separating the introduced constraints: a MIP formulation and an enumeration algorithm. We also show how to take these constraints into account in the pricing problem, which turns into the knapsack problem with hard and soft conflicts (KPHSC), a generalization of the KPC. We formulate the KPHSC as a mixed integer program (MIP). Finally, for the two special cases of the BPPTL mentioned, we propose an exact branch-cut-and-price (BCP) algorithm which uses our cut separation routines, a branch-and-bound algorithm for solving the pricing problem, a strong diving primal heuristic, and the Ryan&Foster strong branching. We generate new instances for the BPPTL with $L = \infty$ inspired by the phytosanitary treatments planning application described above. We show that our BCP algorithm significantly outperforms a commercial MIP solver applied to the compact formulation for the problem. We also test our BCP algorithm on literature instances of the BPP-P, the SALBP-1, and the BPP-GP. We are able improve the best known lower and upper bounds for numerous instances. Thus, optimality is reached for the majority of open instances.

The article is organized as follows. In Section 2 we present our formulations of BPPTL. A BCP algorithm for the case $L = \infty$ is presented in Section 3. A modification of this algorithm for the case with $L = 1$ and non-negative time lags is given in Section 4. Computational results are presented in Section 5. We draw conclusions and discuss future work in Section 6.

2 Integer programming formulations

2.1 A compact formulation

Recall that L is the maximum number of bins per time period. Let $T \in \mathbb{Z}_+$ be an upper bound to the number of time periods required to assign all the bins in an optimal solution, $\mathcal{L} = \{1, \dots, L\}$, and $\mathcal{T} = \{1, \dots, T\}$. Let binary variable $x_{i,b,p}$ take value one if and only if item $i \in V$ is assigned to bin (b, p) with $b \in \mathcal{L}$ and $p \in \mathcal{T}$. Also let binary variable $u_{b,p}$ take value one if only if bin (b, p) with $b \in \mathcal{L}$ and $p \in \mathcal{T}$ has at least one item assigned to it. The BPPTL can then be formulated as the following IP.

$$\min \sum_{b\in\mathcal{L}}\sum_{p\in\mathcal{T}} u_{b,p} \tag{2a}$$

$$\text{s.t.} \sum_{b\in\mathcal{L}}\sum_{p\in\mathcal{T}} x_{i,b,p} = 1 \qquad \forall i \in V, \tag{2b}$$

$$\sum_{i\in V} w_i x_{i,b,p} \le W u_{b,p} \qquad \forall b \in \mathcal{L}, p \in \mathcal{T}, \tag{2c}$$

$$l_{i,j} + \sum_{p\in\mathcal{T}} p \cdot \sum_{b\in\mathcal{L}} x_{i,b,p} \le \sum_{p\in\mathcal{T}} p \cdot \sum_{b\in\mathcal{L}} x_{j,b,p} \qquad \forall (i,j) \in A, \tag{2d}$$

$$x_{i,b,p} \in \{0,1\} \qquad \forall i \in V, b \in \mathcal{L}, p \in \mathcal{T}, \tag{2e}$$

$$u_{b,p} \in \{0,1\} \qquad \forall b \in \mathcal{J}, p \in \mathcal{T}. \tag{2f}$$

In this formulation, constraints (2b) require each item to be assigned to exactly one bin. Constraints (2c) impose two conditions: i) if an item is assigned to a bin (b,p) then the corresponding variable $u_{b,p}$ is equal to one; ii) the items assigned to a bin must have a cumulative weight smaller than or equal to the capacity of the bin. Constraints (2d) guarantee that the time lags are satisfied. This formulation can be improved by disaggregating constraints (2c), or breaking some symmetries. We discuss these variants in Appendix B in the online supplement (https://pubsonline.informs.org/doi/suppl/10.1287/ijoc.2022.1165).

Given a directed path in graph G from node i to node j, its length is the sum of the lags for the arcs in the path. Let $d(i,j)$ represent the length of the longest path from node i to node j in G. We convene $d(i,j) = -\infty$ if there is no path from i to j.

For each $i \in V$, let $es(i)$ and $ls(i)$ respectively represent the earliest and latest time periods in which i can be assigned to a bin. A natural pre-processing consists in fixing to zero all variables $x_{i,b,p}$ such that $i \in V$, $b \in \mathcal{L}$, and $p < es(i)$ or $p > ls(i)$. An approach to compute the values $es(i)$ and $ls(i)$ consists in adding two extra nodes to the graph G: the source node s fixed to artificial time period 0 and the sink node f fixed to artificial time period $T+1$. Then for each item $i \in V$ we add arcs (s,i) and (i,f) with time lag one. Then $es(i) = d(s,i)$ and $ls(i) = T + 1 - d(i,f)$.

2.2 A Dantzig-Wolfe reformulation

A natural Dantzig-Wolfe reformulation of model (2) consists in convexifying constraints (2c). Given $p \in \mathcal{T}$, let $\mathcal{B}_p$ be the collection of sets of items which may be assigned to a bin in time period p: $B \in \mathcal{B}_p$ if and only if $\sum_{i\in B} w_i \le W$, and $es(i) \le p \le ls(i)$, $\forall i \in B$. For each $B \in \mathcal{B}_p$, we define a binary variable λ_B taking value one if and only if set B of items occupies one bin in period p in the solution. We denote by $\mathbb{1}_B$ the indicator function of B: $\mathbb{1}_B(i) = 1$ if $i \in B$, and $\mathbb{1}_B(i) = 0$ otherwise, for all $i \in V$. The reformulation is then the following.

$$\min \sum_{p\in\mathcal{T}}\sum_{B\in\mathcal{B}_p} \lambda_B \tag{3a}$$

$$\text{s.t.} \sum_{p\in\mathcal{T}}\sum_{B\in\mathcal{B}_p} \mathbb{1}_B(i)\lambda_B = 1 \qquad \forall i \in V, \tag{3b}$$

$$l_{i,j} + \sum_{p\in\mathcal{T}} p \cdot \sum_{B\in\mathcal{B}_p} \mathbb{1}_B(i)\lambda_B \le \sum_{p\in\mathcal{T}} p \cdot \sum_{B\in\mathcal{B}_p} \mathbb{1}_B(j)\lambda_B \qquad \forall (i,j) \in A, \tag{3c}$$

$$\sum_{B\in\mathcal{B}_p} \lambda_B \le |\mathcal{L}| \qquad \forall p \in \mathcal{T}, \tag{3d}$$

$$\lambda_B \in \{0,1\} \qquad \forall p \in \mathcal{T}, B \in \mathcal{B}_p. \tag{3e}$$

The linear relaxation of model (3) can be solved by the column generation approach in which the pricing problem is the standard $0-1$ knapsack problem. It is known however (Pereira,

2016) that precedence constraints (3c) do not affect the value of the lower bound obtained by such relaxation for the BPP-P and the SALBP-1. We also show below in our computational experiments that the branch-and-price algorithm that solves formulation (3) does not have a good performance because of a bad quality of the lower bound given by the linear relaxation.

Thus in the following section, we propose an alternative extended formulation for the BPPTL, which provides a better relaxation. Moreover, it allows us to exploit the symmetry related to time periods. Then, in the following two sections we will show how to turn this relaxation into an exact formulation for two important special cases of the problem. This is done by defining an exponential family of constraints that cut off infeasible solutions of the relaxation.

2.3 An alternative relaxation

In any feasible solution of the BPPTL, a set B of items assigned to the same bin must satisfy the following conditions:

$$\sum_{i \in B} w_i \leq W \tag{4a}$$

$$d(i,j) \leq 0 \text{ and } d(j,i) \leq 0 \qquad \forall i,j \in B. \tag{4b}$$

Constraint (4a) follows from (1a) and constraints (4b) follow from (1b). These two conditions state that the bin capacity constraint has to be satisfied, and that (transitive) lag constraints forbid some items from being packed in the same bin.

Let $\mathcal{B}$ be the collection of all possible non-empty sets B of items satisfying constraints (4a) and (4b). For each $B \in \mathcal{B}$, we define a binary variable λ_B taking value one if and only if set B of items occupies one bin in the solution. A relaxation of the BPPTL then can be formulated as follows.

$$\min \sum_{B \in \mathcal{B}} \lambda_B \tag{5a}$$

$$\text{s.t.} \sum_{B \in \mathcal{B}} \mathbb{1}_B(i) \lambda_B = 1 \qquad \forall i \in V, \tag{5b}$$

$$\lambda_B \in \{0,1\} \qquad \forall B \in \mathcal{B}. \tag{5c}$$

In this formulation the objective function is to minimize the number of bins selected in the solution. This relaxation has a similar structure to the one used for the simple assembly line balancing problem (Pereira, 2015) and the bin packing problem with precedences (Pereira, 2016), albeit with a different definition of $\mathcal{B}$. Moreover, integrality constraints (5c) are relaxed to the linear constraints $0 \leq \lambda_B \leq 1$ $\forall B \in \mathcal{B}$ in (Pereira, 2015, 2016).

A feasible solution $\bar{\lambda}$ for model (5) does not necessarily define a feasible solution for the BPPTL, since it is not always possible to assign bins in set $\{B \in \mathcal{B} : \bar{\lambda}_B = 1\}$ to time periods while satisfying the time lag constraints. This can be seen in the example illustrated in Figure 3. In this example, a feasible solution to (5) has items $\{A, B\}$ assigned to one bin and items $\{C, D\}$ to another. However, it is not possible to assign those bins to time periods satisfying the time lag constraints. In fact, item A must be assigned before D, and item C must be assigned before B. One can see that the issue stems from the fact that when items are assigned to bins, lags between items become lags between bins, and cycles of positive length may appear. This concept of cycles between bins is formalized in Section 3, and is used to characterize infeasible solutions.

3 The case with unlimited number of available bins

In this section we assume that $L = \infty$, i.e., there is no restriction on the number of bins that can be assigned to a time period. We denote this variant of the problem as the BPPTL$^{\infty}$.

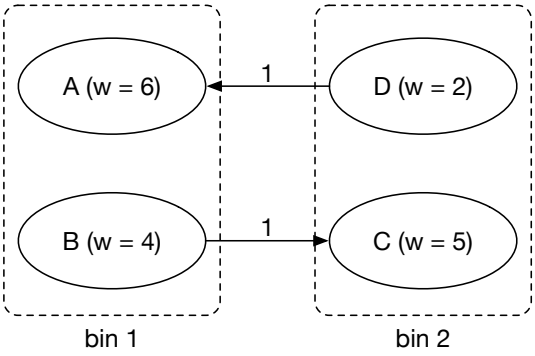


Figure 3: Example of instance in which a feasible solution to relaxation (5) does not define a feasible solution for the BPPTL. Here $W = 10$.

We propose in this section a branch-cut-and-price algorithm based on relaxation (5) to solve this variant to optimality. In Section 3.1, we give a characterization of feasible solutions of the BPPTL$^\infty$ expressed by constraints involving only variables λ_B. Later, in Section 3.2, we describe the components of the branch-cut-and-price algorithm, i.e., separation algorithms for the cuts, an IP formulation and a dedicated branch-and-bound for the pricing problem, primal heuristic and branching.

3.1 Characterization of valid assignments

Let $\mathcal{P}$ be a partition of V. We denote by $\mathcal{P}(i)$ the element $B \in \mathcal{P}$ such that $i \in B$. For each $B, B' \in \mathcal{P}$, we define $\bar{d}(B, B') = \max\{d(i,j) : i \in B, j \in B'\}$ as the maximum distance (in terms of time lags) between an item in B and an item in B'. We introduce the notion of *aggregated graph*, which is useful for defining the necessary and sufficient conditions for a feasible solution to (5) to be feasible for the BPPTL.

Definition 1. *Let $G = (V, A, l)$ be a directed valued graph, and $\mathcal{P}$ be a partition of V. The* aggregated graph *of G induced by $\mathcal{P}$ is the valued digraph $G^{\mathcal{P}} = (\mathcal{P}, A^{\mathcal{P}}, \bar{d})$, where $\mathcal{P}$ is the set of nodes in the graph, $A^{\mathcal{P}} \coloneqq \left\{(B, B') \subseteq \mathcal{P} \times \mathcal{P} : \bar{d}(B, B') > -\infty\right\}$ is the set of arcs in the graph, and $\bar{d}(B, B')$ is the length of arc $(B, B') \in A^{\mathcal{P}}$.*

We now show that whenever the aggregated graph induced by a partition is acyclic, the elements of the partition can be assigned to time periods satisfying the time lag constraints.

Proposition 1. *Let $G = (V, A, l)$ be a directed valued graph and $\mathcal{P}$ be a partition of V. Also let $G^{\mathcal{P}} = (\mathcal{P}, A^{\mathcal{P}}, \bar{d})$ be the aggregated graph of G induced by $\mathcal{P}$. There is a mapping $\tau : \mathcal{P} \longrightarrow \mathbb{Z}_+$, such that for each $(i,j) \in A$, $\tau(\mathcal{P}(i)) + l_{i,j} \leq \tau(\mathcal{P}(j))$ if and only if there is no cycle of positive length in $G^{\mathcal{P}}$.*

Proof. The proof relies on the fact that graph G has no cycle of positive length if and only if a mapping $\tau : V \longrightarrow \mathbb{Z}_+$ exists, such that for each $(i,h) \in A$, $\tau(i) + l_{i,h} \leq \tau(j)$ (see Lemma 1 in Appendix A in the online supplement, https://pubsonline.informs.org/doi/suppl/10.1287/ijoc.2022.1165). Consider this mapping τ for the graph $G^{\mathcal{P}}$. Then for each $(B, B') \in A^{\mathcal{P}}$, it holds that $\tau(B) + \bar{d}(B, B') \leq \tau(B')$. But for each $i, j \in A$, we have $l_{i,j} \leq d(i,j)$, and since $\bar{d}(B, B') = \max\{d(i,j) : i \in B, j \in B'\}$, then $d(i,j) \leq \bar{d}(\mathcal{P}(i), \mathcal{P}(j))$. This implies that for each $(i,j) \in A$, $\tau(\mathcal{P}(i)) + l_{i,j} \leq \tau(\mathcal{P}(i)) + \bar{d}(\mathcal{P}(i), \mathcal{P}(j)) \leq \tau(\mathcal{P}(j))$. □

We call a partition $\mathcal{P} \subseteq \mathcal{B}$ *suitable* if graph $G^{\mathcal{P}}$ does not have a directed cycle of positive length, and we say that $\mathcal{P}$ is *unsuitable* otherwise. We denote by $\mathcal{N}$ the family of *unsuitable* partitions of V.

The partition of items induced by a feasible solution of an instance of the BPPTL$^\infty$ must be *suitable*, because the assignment of bins to time periods implies that there are no positive cycles in the aggregated graph. But Proposition 1 also shows that any *suitable* partition $\mathcal{P} \subseteq \mathcal{B}$ induces a feasible solution for the instance, and the proof also gives a procedure to recover the actual solution (the time of each bin) from the partition by computing the longest distance from an artificial source node to each node in the aggregated graph, which can be done in time $\mathcal{O}\left(|\mathcal{P}|\,|A^{\mathcal{P}}|\right)$.

A classical way to ensure feasibility is to add the following set of constraints (known as *no-good cuts*) to formulation (5):

$$\sum_{B \in \mathcal{P}} \lambda_B \leq |\mathcal{P}| - 1 \qquad \forall \mathcal{P} \in \mathcal{N}. \tag{6}$$

One of our contributions is to avoid using these constraints by finding a better characterization of feasible solutions.

Formulation (5) together with constraints (6) defines an exact formulation of the BPPTL$^\infty$. However, each constraint in (6) forbids only one *unsuitable* partition, and the coefficient of a variable λ_B in these constraints depends on set B. This makes dynamic generation of variables λ very difficult. We now introduce an alternative to constraints (6), which defines a computationally tractable equivalent formulation.

The alternative constraints rely on the fact that cycles in the aggregated graphs are related to lags between pairs of items. Consider a partition $\mathcal{P}$ such that the induced aggregated graph $G^{\mathcal{P}}$ has a cycle of positive length $(B_1, B_2, \ldots, B_R)$, and denote $B_{R+1} = B_1$. For each arc (B_r, B_{r+1}) in the cycle there are elements $t_r \in B_r$ and $h_{r+1} \in B_{r+1}$ such that $\bar{d}(B_r, B_{r+1}) = d(t_r, h_{r+1})$. If we consider the set $\{(h_1, t_1), \ldots, (h_R, t_R)\}$ of pairs of items, any partition where each of these pairs is in the same bin induces an aggregated graph with a positive length cycle. Here notations t and h stand for the "tail" and "head" of the corresponding arc.

Figure 4 presents an example of a partition that induces a positive cycle in the aggregated graph. Any partition in which pairs of nodes (h_1, t_1), (h_2, t_2) and (h_3, t_3) each belong to the same block of the partition, induces the aggregated graph with a positive length cycle.

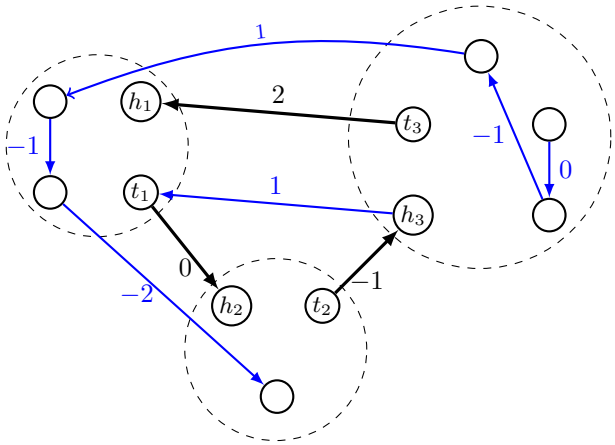


Figure 4: Example of partition that induces a positive cycle. Here $(h_1, h_2, h_3, t_1, t_2, t_3) \in \mathcal{C}$.

We now define the collection of all sets of pairs which induce a positive length cycle in the aggregated graph, and thus an infeasible solution.

Definition 2. *For each $R \in \mathbb{N}$, $R \geq 2$, we define $\mathcal{C}_R \subseteq V^{2R}$ as the collection of all tuples $(h_1, \ldots, h_R, t_1, \ldots, t_R)$ that satisfy the following properties.*

- *All elements $h_1, \ldots, h_R, t_1, \ldots t_R$ are different.*

- *For each* $r \in \{1, \dots, R\}$, $d(h_r, t_r) \leq 0$ *and* $d(t_r, h_r) \leq 0$.
- $\sum_{r=1}^{R} d(t_r, h_{r+1}) \geq 1$, *where* $h_{R+1} = h_1$.

We also denote $\mathcal{C} = \bigcup_{R \geq 2} \mathcal{C}_R$, and for each tuple $C = (h_1, \dots, h_R, t_1, \dots, t_R) \in \mathcal{C}$ we define as $\mathcal{F}_C$ the set of pairs of items which constitute C: $\mathcal{F}_C = \big\{\{h_1, t_1\}, \{h_2, t_2\}, \dots, \{h_R, t_R\}\big\}$.

We can now characterize an *unsuitable* partition as a partition that contains a forbidden set of pairs.

Proposition 2. *A partition* $\mathcal{P} \subseteq \mathcal{B}$ *is* unsuitable *if and only if* $C \in \mathcal{C}$ *exists, such that for each pair* $\{h_r, t_r\} \in \mathcal{F}_C$, $\mathcal{P}$ *contains an element* B_r *containing both* h_r *and* t_r*:* $\{h_r, t_r\} \subseteq B_r \in \mathcal{P}$.

Proof. Let $\mathcal{P} \subseteq \mathcal{B}$ be a partition of V, and let $C \in \mathcal{C}_R$, $R \geq 2$, be such that each $\{h_r, t_r\} \in \mathcal{F}_C$ is contained in $B_r \in \mathcal{P}$ for each $r \in \{1, \dots, R\}$. Then, by definition of $\bar{d}$, for each $r \in \{1, \dots, R\}$ we have $\bar{d}(B_r, B_{r+1}) \geq d(t_r, h_{r+1})$, where $B_{R+1} = B_1$. This implies that $\sum_{r=1}^{R} \bar{d}(B_r, B_{r+1}) \geq 1$, and then there is a circuit in $(G^{\mathcal{P}}, \bar{d})$ of positive length. Since the circuit can be decomposed into multiple cycles, and the length of the circuit is equal to the sum of the length of those cycles, one of those cycles must be of positive length, and then $\mathcal{P}$ must be *unsuitable*.

Assume now that $\mathcal{P}$ is *unsuitable*, and consider a cycle of positive length in graph $G^{\mathcal{P}}$ with the minimum possible number of nodes. Let R be the number of nodes of that cycle, and let the nodes of the cycle be $B_1, B_2, \dots, B_R$.

For each arc (B_r, B_{r+1}) in this cycle, let $t_r \in B_r$ and $h_{r+1} \in B_{r+1}$ be such that $d(t_r, h_{r+1}) = \bar{d}(B_r, B_{r+1})$, where $B_{R+1} = B_1$. Thus we have $\{h_r, t_r\} \subseteq B_r$, and since the cycle in the aggregated graph has positive length, we have $\sum_{r=1}^{R} d(t_r, h_{r+1}) = \sum_{r=1}^{R} \bar{d}(B_r, B_{r+1}) \geq 1$.

We now show that all these elements $(h_1, \dots, t_R)$ are different. Since all B_r are different elements of a partition of V, for each $r, r' \in \{1, \dots, R\}$ with $r \neq r'$ it holds $h_r \neq h_{r'}$, $t_r \neq t_{r'}$, and $h_r \neq t_{r'}$. If $h_r = t_r$ then $\bar{d}(B_{r-1}, B_{r+1}) \geq d(t_{r-1}, h_r) + d(t_r, h_{r+1}) = \bar{d}(B_{r-1}, B_r) + \bar{d}(B_r, B_{r+1})$, and then removing B_r from the cycle would still give a cycle of positive length, which contradicts the minimality of the cycle considered. Then, all the elements $h_1, \dots, t_R$ are different.

Finally, since $B_r \in \mathcal{B}$, then $d(h_r, t_r) \leq 0$ and $d(t_r, h_r) \leq 0$. □

The following reformulation is valid for the BPPTL$^\infty$ due to Proposition 2.

$$\min \sum_{B \in \mathcal{B}} \lambda_B \tag{7a}$$

$$\text{s.t.} \sum_{B \in \mathcal{B}} \mathbb{1}_B(i) \lambda_B = 1 \qquad \forall i \in V, \tag{7b}$$

$$\sum_{r=1}^{R} \sum_{\substack{B \in \mathcal{B}, \\ \{h_r, t_r\} \subseteq B}} \lambda_B \leq R - 1 \qquad \forall R \geq 2, \forall (h_1, \dots, t_R) \in \mathcal{C}_R, \tag{7c}$$

$$\lambda_B \in \{0, 1\} \qquad \forall B \in \mathcal{B}. \tag{7d}$$

3.2 The branch-cut-and-price algorithm

We develop a branch-cut-and-price approach to solve formulation (7). At each node of the branch-and-bound tree the linear relaxation is solved by the following column and cut generation procedure.

- *Step 0:* Start with an initial set of variables λ and constraints (7c).

- *Step 1:* Solve the restricted master problem (RMP). The RMP is the linear relaxation of (7) with the current restricted subset of columns and constraints.
- *Step 2:* Solve the pricing problem which looks for a variable λ_B with $B \in \mathcal{B}$ with negative reduced cost. If such a column is found, add λ_B to the RMP and go back to *Step 1.*
- *Step 3:* Solve the separation problem. The separation problem consists in finding a constraint (7c) violated by the current solution of the RMP.

Approaches for solving the separation problem are presented in Section 3.2.1. The MIP formulation for the pricing problem is proposed in Section 3.2.2. Other components of the BCP algorithm are described in Section 3.2.4.

3.2.1 Separation.

Given a solution $\bar{\lambda} = \left(\bar{\lambda}_B\right)_{B\in\mathcal{B}}$ that satisfies (7b), we are looking for a constraint in (7c) that is violated by $\bar{\lambda}$.

If $\bar{\lambda}$ satisfies (7d), i.e., it is an integer solution, it induces a partition $\mathcal{P}$ of the items, and finding a violated constraint in (7c) is equivalent to finding a cycle of positive length in the aggregated graph $G^{\mathcal{P}}$, according to Proposition 2. We first construct the aggregated graph $G^{\mathcal{P}}$. Then we compute the longest distance between each pair of nodes in this graph, which can be done in time $\mathcal{O}\left(|\mathcal{P}|^3\right)$ with the Floyd-Warshall algorithm. Finally, we look for a cycle of a positive length by checking if the distance from any node to itself is positive. In the event that a cycle of positive length is found, the proof of Proposition 2 describes the procedure to find the tuple $(h_1, \ldots, t_R) \in \mathcal{C}$ and to generate the associated violated constraint.

Separating integer solutions is enough for the algorithm to be valid. However, in order to strengthen the relaxation, we also propose procedures to separate fractional solutions. To do so, we define an auxiliary directed multigraph in which finding a cycle with some specific properties is equivalent to finding a violated constraint in (7c).

Let $\left(\bar{\lambda}_B\right)_{B\in\mathcal{B}}$ be a solution satisfying (7b) but not necessarily (7d). For each $i, j \in V$ with $i \neq j$ we define the value $\xi_{i,j} = \sum_{\substack{B\in\mathcal{B},\\ \{i,j\}\subseteq B}} \bar{\lambda}_B$. Then we construct multigraph $\mathcal{G} = (V, \mathcal{A})$ with two labels in each arc. Each arc $a \in \mathcal{A}$ is represented by a tuple $(h(a), t(a), l(a), v(a)) \in V \times V \times \mathbb{Z} \times \mathbb{R}$, where $h(a)$ is the head of the arc, and $t(a)$ is the tail of the arc. The set of arcs is composed of two subsets, $\mathcal{A} = \mathcal{A}_G \cup \mathcal{A}_\lambda$, where these subsets are defined as follows:

- For each $(i, j) \in A$, arc a in which $h(a) = i$, $t(a) = j$, $l(a) = d(i, j)$ and $v(a) = 0$ belongs to $\mathcal{A}_G$.
- For each pair $\{i, j\} \subseteq V$ such that $i \neq j$ and $\xi_{i,j} > 0$, arc a in which $h(a) = i$, $t(a) = j$, $l(a) = 0$, and $v(a) = 1 - \xi_{i,j}$ belongs to $\mathcal{A}_\lambda$.

There is one arc in $\mathcal{A}_G$ for each arc in the original graph G, and two arcs in $\mathcal{A}_\lambda$, one in each direction, for each pair of nodes that are contained in any $B \in \mathcal{B}$ such that $\bar{\lambda}_B > 0$. Note that in this graph there might be two arcs with the same head and tail, one of them in $\mathcal{A}_G$ and the other in $\mathcal{A}_\lambda$.

Proposition 3. *Let* $\left(\bar{\lambda}_B\right)_{B\in\mathcal{B}}$ *be a solution satisfying* (7b)*. Then it satisfies all constraints in* (7c) *if and only if multigraph* $\mathcal{G} = (V, \mathcal{A})$ *constructed from* $\bar{\lambda}$ *has no cycle* $(a_1, a_2, \ldots, a_K)$ *such that* $\sum_{k=1}^{K} l(a_k) \geq 1$ *and* $\sum_{k=1}^{K} v(a_k) < 1$.

Proof. ($\Rightarrow$) Assume first that $\bar{\lambda}$ violates a constraint in (7c), then there is $R \geq 2$ and $(h_1, \ldots, t_R) \in$

$\mathcal{C}_R$ such that

$$\sum_{r=1}^{R} \sum_{\substack{B\in\mathcal{B},\\ \{h_r,t_r\}\subseteq B}} \bar{\lambda}_B > R-1.$$

This is equivalent to

$$1 > \sum_{r=1}^{R} \left(1 - \sum_{\substack{B\in\mathcal{B},\\ \{h_r,t_r\}\subseteq B}} \bar{\lambda}_B\right) = \sum_{r=1}^{R} (1-\xi_{h_r,t_r}).$$

With this we define the following cycle with $K=2R$ arcs in graph $\mathcal{G}$. For each $r\in\{1,\dots,R\}$, we take arc $a_{2r-1}\in\mathcal{A}_\lambda$ such that $t(a_{2r-1})=h_r$ and $h(a_{2r-1})=t_r$, and we take arc $a_{2r}\in\mathcal{A}_G$ such that $t(a_{2r})=t_r$ and $h(a_{2r})=h_{r+1}$. We have $v(a_{2r-1})=1-\xi_{h_r,t_r}$ and $v(a_{2r})=0$. Summing over all arcs in the cycle, we obtain $\sum_{k=1}^{2R} v(a_k)<1$. We also have $l(a_{2r-1})=0$ and $l(a_{2r})=d(t_r,h_{r+1})$. Then by definition of $\mathcal{C}_R$ we have $\sum_{k=1}^{2R} l(a_k)=\sum_{r=1}^{R} d(t_r,h_{r+1})\geq 1$.

($\Leftarrow$) Assume now that there are cycles $(a_1,a_2,\dots,a_K)$ in graph $\mathcal{G}$ satisfying both $\sum_{k=1}^{K} l(a_k)\geq 1$ and $\sum_{k=1}^{K} v(a_k)<1$. Among these cycles we take one with the minimum number K of arcs. We now show that the arcs in this cycle are alternating between arcs of $\mathcal{A}_G$ and $\mathcal{A}_\lambda$. Let a and a' be two consecutive arcs in the cycle, with $h(a)=i$, $t(a)=j$, $h(a')=j$ and $t(a')=k$. Now consider two cases.

i) If both arcs belong to $\mathcal{A}_G$, then we can replace both arcs by arc (i,k) in $\mathcal{A}_G$ since $d(i,k)\geq d(i,j)+d(j,k)$.

ii) If both arcs belong to $\mathcal{A}_\lambda$, then we have

$$\begin{aligned}
\xi_{i,j}+\xi_{j,k} &= \sum_{\substack{B\in\mathcal{B},\\ \{i,j\}\subseteq B}} \bar{\lambda}_B + \sum_{\substack{B\in\mathcal{B},\\ \{j,k\}\subseteq B}} \bar{\lambda}_B \\
&= \underbrace{\sum_{\substack{B\in\mathcal{B},\\ \{i,j\}\subseteq B,k\notin B}} \bar{\lambda}_B + \sum_{\substack{B\in\mathcal{B},\\ \{i,j\}\subseteq B,k\in B}} \bar{\lambda}_B + \sum_{\substack{B\in\mathcal{B},\\ \{j,k\}\subseteq B,i\notin B}} \bar{\lambda}_B}_{\leq \sum_{\substack{B\in\mathcal{B},\\ j\in B}} \bar{\lambda}_B} + \underbrace{\sum_{\substack{B\in\mathcal{B},\\ \{j,k\}\subseteq B,i\in B}} \bar{\lambda}_B}_{\leq \sum_{\substack{B\in\mathcal{B},\\ \{i,k\}\subseteq B}} \bar{\lambda}_B} \\
&\leq 1+\xi_{i,k}.
\end{aligned}$$

The first consequence of this inequality is that $\xi_{i,k}>0$, and arc (i,k) belongs to $\mathcal{A}_\lambda$, as otherwise we would have $v(a)+v(a')\geq 1$. The second consequence is that $v(a)+v(a')=2-\xi_{i,j}-\xi_{j,k}\geq 1-\xi_{i,k}$, and value v for arc (i,k) is not larger than $v(a)+v(a')$. Then we can replace arcs a and a' by arc (i,k).

In both cases, there is a contradiction with the minimality of the number of arcs in the cycle. Thus we can conclude that the arcs are alternating between $\mathcal{A}_G$ and $\mathcal{A}_\lambda$. Thus K must be even, and we can denote $K=2R$. Without loss of generality we assume $a_1\in\mathcal{A}_\lambda$. We define $h_r=t(a_{2r-1})$ and $t_r=h(a_{2r-1})$ for each $r\in\{1,\dots,R\}$. Now we are going to show that $(h_1,\dots,t_R)\in\mathcal{C}_R$.

The minimality of the cycle implies that all nodes $(h_1,\dots,t_R)$ are different. Since arc (h_r,t_r) belongs to $\mathcal{A}_\lambda$, there must be some $B\in\mathcal{B}$ with $\bar{\lambda}_B>0$ such that $\{h_r,t_r\}\subseteq B$. By definition of $\mathcal{B}$, this implies that $d(h_r,t_r)\leq 0$ and $d(t_r,h_r)\leq 0$. Finally, since $a_{2r-1}\in\mathcal{A}_\lambda$, $l(a_{2r-1})=0$, then $\sum_{r=1}^{R} l(a_{2r})\geq 1$, which implies that $\sum_{r=1}^{R} d(t_r,h_{r+1})\geq 1$. □

The first approach to find such a cycle in multigraph $\mathcal{G}=(V,\mathcal{A})$ is based on an IP. Let y_a be a binary variable which determines whether arc $a\in\mathcal{A}$ participates in the cycle or not. Consider

the following IP formulation.

$$\min \sum_{a\in\mathcal{A}} v(a)y_a \tag{8a}$$

$$\text{s.t.} \sum_{a\in\mathcal{A}} l(a)y_a \geq 1, \tag{8b}$$

$$\sum_{a\in\mathcal{A}\,:\,h(a)=i} y_a = \sum_{a\in\mathcal{A}\,:\,t(a)=i} y_a \qquad \forall i \in V, \tag{8c}$$

$$y_a \in \{0,1\} \qquad \forall a \in \mathcal{A}. \tag{8d}$$

It can be seen that there is a cycle $(a_1, a_2, \ldots, a_K)$ in $\mathcal{G}$ such that $\sum_{k=1}^{K} l(a_k) \geq 1$ and $\sum_{k=1}^{K} v(a_k) < 1$ if and only if the optimal value of (8) is strictly smaller than one. Constraints (8c) force the solution to be a collection of cycles. Due to constraints (8b), at least one cycle in this collection must have a positive sum of values l. If the objective function is strictly less than one, then any cycle in the collection has the sum of values v strictly less than one.

Generating several constraints in each separation round helps to improve convergence of the cut generation procedure. This can be achieved by solving formulation (8) multiple times. After each solving, a constraint is added to avoid selecting the same cut again. Suppose that cycle $(a_1, a_2, \ldots, a_K)$ is found after solving the IP. Without loss of generality, let $a_1 \in \mathcal{A}_G$ and $R = 2K$. Then the constraint $\sum_{r=1}^{R} y_{a_{2r}} \leq R - 1$ is added to (8), which is resolved again.

We now describe the second approach to find a desired cycle in multigraph $\mathcal{G} = (V, \mathcal{A})$. Given a set A' of arcs, we denote $v(A') = \sum_{a\in A'} v(a)$ and $l(A') = \sum_{a\in A'} l(a)$, and given two nodes $i, j \in V$ we denote $v_{i,j}$ the minimum possible sum of v values of a path from i to j. The approach consists in a partial enumeration of cycles C in $\mathcal{G}$ satisfying $v(C) < 1$ and $l(C) \geq 1$.

We perform one iteration for each node $i \in V$. In this iteration, we try to find cycles C containing node i such that $v(C) < 1$ and $l(C) \geq 1$. Then we mark node i to exclude cycles containing i from the search in subsequent iterations. In each iteration, we use a recursive procedure which takes the current partial path from node i to current node j, containing set C of arcs and set $V_C \cup \{j\}$ of nodes. This procedure iterates over arcs $a \in \delta^+(j)$. If head $h(a)$ of arc a has been marked, arc a is ignored. Otherwise, if node $h(a)$ is contained in V_C then a subset C' of the arcs in $C \cup \{a\}$ forms a cycle. In this case, if conditions $v(C') < 1$ and $l(C') \geq 1$ held, then we store cycle C' inducing a violated inequality. If node $h(a)$ is neither blocked nor contained in V_C we check conditions $v(C) + v(a) + v(a') < 1$ and $l(C) + l(a) + l(a') \geq 1$, where $a' = (h(a), i)$. If none of these conditions is true, it is unlikely that arcs in $C \cup \{a\}$ are contained in a cycle inducing a valid inequality, and arc a is ignored. If both conditions are true, we augment the partial path with arc a and recursively call the same procedure. The separation algorithm by partial enumeration is formally presented in Algorithm 1.

3.2.2 Pricing problem.

Consider an optimal dual solution $(\bar{\pi}, \bar{\mu})$ of the restricted master problem, where $(\pi_i)_{i\in V}$ are the dual variables of constraints (7b) and $(\mu_C)_{C\in\mathcal{C}}$ are the dual variables of constraints (7c).

The pricing problem consists in finding a set of items $B \in \mathcal{B}$ such that the reduced cost of variable λ_B is minimized. The coefficient of λ_B in the constraint (7b) corresponding to $i \in V$ is one if $i \in B$, and zero otherwise. The coefficient of λ_B in the constraint (7c) corresponding to tuple $C \in \mathcal{C}$ is equal to the number of item pairs $\{i, j\} \in \mathcal{F}_C$ such that both i and j are contained in B. We define $\theta(B, C) = |\{\{i, j\} \in \mathcal{F}_C : i \in B, j \in B\}|$. Then the reduced cost of variable λ_B is $1 - \sum_{i\in B} \bar{\pi}_i - \sum_{C\in\mathcal{C}} \theta(B, C)\bar{\mu}_C$.

We now describe a MIP to find a set B with minimum reduced cost. Let $\bar{\mathcal{C}}$ be the set of active constraints (7c): $\bar{\mathcal{C}} = \{C \in \mathcal{C} : \bar{\mu}_C < 0\}$. Also let $\bar{\mathcal{F}} = \bigcup_{C\in\bar{\mathcal{C}}} \mathcal{F}_C$. Let binary variable x_i be equal

Function `DFSRecursion`(i, j, V_C, C)**:**
  for $a \in \delta^+(j)$ **do**
    if $h(a)$ *is marked* **then continue**
    if $h(a) \in V_C$ **then**
      Find $C' \subseteq C \cup \{a\}$ that forms a cycle.
      if $v(C') < 1$ **and** $l(C') \geq 1$ **then** store cycle C'
      continue
    end
    $a' \longleftarrow (h(a), i)$
    if $v(C) + v(a) + v(a') \geq 1$ **or** $l(C) + l(a) + l(a') \leq 0$ **then continue**
    `DFSRecursion`(i, $h(a)$, $V_C \cup \{j\}$, $C \cup \{a\}$)
  end
end

Function `Main`($\mathcal{G}$, v, l)**:**
  for $i \in V$ **do**
    `DFSRecursion`(i, i, $\emptyset$, $\emptyset$)
    Mark node i
  end
end

Algorithm 1: Fractional separation by partial enumeration

to one if item $i \in V$ belongs to B, and zero otherwise. To simplify the notation, we introduce $\psi_{i,j} = \sum_{C \in \bar{\mathcal{C}}: \{i,j\} \in \mathcal{F}_C} \bar{\mu}_C$. Also let binary variable $y_{i,j}$ be equal to one if both items from pair $\{i,j\} \in \bar{\mathcal{F}}$ belong to B, and zero otherwise. Then the pricing problem can be formulated as the following MIP.

$$\max \sum_{i \in V} \bar{\pi}_i x_i + \sum_{\{i,j\} \in \bar{\mathcal{F}}} \psi_{i,j} y_{i,j} \tag{9a}$$

$$\text{s.t.} \sum_{i \in V} w_i x_i \leq W, \tag{9b}$$

$$x_i + x_j \leq 1 \qquad \forall \{i,j\} \subseteq V, d(i,j) \geq 1 \text{ or } d(j,i) \geq 1, \tag{9c}$$

$$x_i + x_j \leq 1 + y_{i,j} \qquad \forall \{i,j\} \in \bar{\mathcal{F}}, \tag{9d}$$

$$x_i \in \{0,1\} \qquad \forall i \in V, \tag{9e}$$

$$y_{i,j} \geq 0 \qquad \forall \{i,j\} \in \bar{\mathcal{F}}. \tag{9f}$$

Given an optimal solution $(\bar{x}, \bar{y})$ to (9), we add to the restricted master problem variable $\lambda_{\bar{B}}$, where $\bar{B} = \{i \in V : \bar{x}_i = 1\}$. Constraints (9b) and (9c) require that $\bar{B} \in \mathcal{B}$. Since $\psi_{i,j} < 0$ for $\{i,j\} \in \bar{\mathcal{F}}$ and we are maximizing, each variable $\bar{y}_{i,j}$ takes the minimum possible value. Therefore, $\bar{y}_{i,j}$ is equal to one if and only if both $i \in \bar{B}$ and $j \in \bar{B}$. Therefore,

$$\sum_{\{i,j\} \in \bar{\mathcal{F}}} \sum_{\substack{C \in \bar{\mathcal{C}}: \\ \{i,j\} \in \mathcal{F}_C}} \mu_C \bar{y}_{i,j} = \sum_{C \in \mathcal{C}} \theta(B,C) \bar{\mu}_C,$$

and the objective value of solution $(\bar{x}, \bar{y})$ is equal to the reduced cost of $\lambda_{\bar{B}}$ with the opposite sign.

The pricing problem can be defined as the knapsack problem with hard and soft conflicts (KPHSC). Hard conflicts correspond to pairs $\{i,j\}$ such that either $d(i,j) \geq 1$ or $d(j,i) \geq 1$. Soft conflicts are pairs that are penalized for being both in the set, and correspond to pairs $\{i,j\} \in \bar{\mathcal{F}}$. The KPHSC generalizes the NP-hard knapsack problem with conflicts. The latter is NP-hard, and it is studied for example in (Sadykov and Vanderbeck, 2013; Bettinelli et al., 2017; Coniglio et al., 2021). To our knowledge, the KPHSC has not yet been studied in the literature.

3.2.3 A branch-and-bound algorithm for the pricing problem.

Solving pricing subproblem directly using a general purpose MIP solver happens to be a bottleneck of the BCP algorithm, especially for large instances. Therefore, we propose a dedicated branch-and-bound algorithm to solve (9) which is inspired from the one by Sadykov and Vanderbeck (2013) for the knapsack problem with conflicts.

Each node of the tree search has a list of selected items V^+, and a list of forbidden items V^-. At each node, the branching scheme consists in choosing an item in $V \setminus (V^+ \cup V^-)$ and adding it to V^+ in the left branch and to V^- in the right branch. A lower (primal) bound is computed at each node by taking into account the profits of the selected items, and the penalties incurred by pairs of selected items. If the value of this bound is better than the value of the best known solution so far, the latter is updated. An upper (dual) bound is obtained as the sum of the lower bound and the overestimation of the maximum profit which can be obtained from items in $V \setminus (V^+ \cup V^-)$. This overestimation is calculated by solving the fractional binary knapsack problem formulated in the following way.

- The knapsack size is equal to the residual bin capacity after packing all items in V^+, i.e., $W - \sum_{i \in V^+} w_i$.
- The profit of an item $i \in V \setminus (V^+ \cup V^-)$ is equal to $\hat{p}_i(V^+) = \bar{\pi}_i + \sum_{j \in V^+} \psi_{i,j}$. This profit takes into account the penalties incurred by packing i with items in V^+. Penalties between pairs of items which are both not in V^+ are disregarded. This is valid since penalties always reduce the value of the solution.

The solution of such fractional binary knapsack problem can be computed in linear time if items are ordered according to a non-decreasing profit/weight ratio $\hat{p}_i(V^+)/w_i$. Thus, the list of items $V \setminus (V^+ \cup V^-)$ is always kept ordered according to this ratio. Each time set V^+ is modified, residual bin capacity, values $\hat{p}_i(V^+)$, the ordering of items, and set V^- (due to hard conflicts) are updated.

3.2.4 Other components of the algorithm.

It is well known that the column generation approach may have convergence issues. To improve convergence, we use the automatic dual price smoothing stabilization technique proposed by Pessoa et al. (2018).

Having a good feasible solution is important for the efficiency of the BCP algorithm. Such a solution provides an upper bound for the optimal objective value, and thus many nodes in the branch-and-bound tree may be pruned by bound. To obtain feasible solutions, we use the strong diving heuristic proposed by Sadykov et al. (2019). We now briefly describe how the standard diving heuristic (Joncour et al., 2010) can be applied for the BPPTL and present its strong diving variant.

After solving a node in the branch-and-bound tree, we have a solution $\bar{\lambda}$ to the linear relaxation of formulation (7). If $\bar{\lambda}$ is integer, the node is pruned, otherwise we apply the diving heuristic. In this algorithm, we select a column λ_B such that its value $\bar{\lambda}_B$ in the fractional solution is the closest to 1. We then add this column to the partial solution and change the right-hand-side values of constraints (7b) corresponding to items in B to zero. Afterwards, the linear relaxation of (7) is resolved by column generation (without cut separation). This iterative process continues until the solution to the linear relaxation becomes integer or until the relaxation becomes infeasible.

In order to increase the efficiency of the diving heuristic, we apply the following problem-specific enhancement to it. Each time a column λ_B is fixed to one, we update lag $l_{i,j}$ to $\max\{l_{i,j}, 0\}$ for every pair $\{i, j\} \subseteq B$. Following this update, we recalculate values $d(i, j)$ and $d(j, i)$, for all pairs $\{i, j\} \subseteq V$, i.e., the lengths of the longest paths between each pair of nodes in the graph. If a value $d(i, j)$ becomes positive, the set of hard conflicts in the pricing problem

is updated, i.e., constraint $x_i + x_j \leq 1$ is added to formulation (9), and columns λ_B such that $\{i, j\} \subseteq B$ are removed from the master problem.

We use the strong diving variant (Sadykov et al., 2019) of this diving heuristic. In this variant, we select up to 10 candidate columns λ_B, and we fix them temporarily to one, one-by-one, and resolve the master problem. Then we select the candidate column which resulted in the smallest increase in the lower bound obtained by solving the linear relaxation of formulation (7).

After applying the strong diving heuristic, if the primal-dual gap is positive, we perform branching. Here we use the Ryan and Foster branching (Ryan and Foster, 1981). To do so, we find a pair of items $\{i, j\} \subseteq V$ such that value $\sum_{B \in \mathcal{B}: \{i,j\} \subseteq B} \bar{\lambda}_B$ is fractional. Then we create two child nodes: in the first we impose the requirement that items i and j are put into the same bin, and in the second that items i and j are assigned to different bins. We respectively add the corresponding constraints to the pricing problem: $x_i = x_j$ in the first branch, and $x_i + x_j \leq 1$ in the second. In the branch-and-bound algorithm, this results in respectively merging items i and j, or adding hard conflict between i and j. We also remove from the master problem columns which do not satisfy newly imposed constraints.

Choosing a pair of items to branch on is an important decision, which has a large impact on the efficiency of the BCP algorithm. To find better branching candidates, we use two-phase strong branching. In the first phase, we take up to 30 branching candidates (i.e., pairs of items), create child nodes for them, and solve only the restricted master problem for them without column and cut generation. Up to three best candidates are then chosen according to the product rule (Achterberg, 2007). For child nodes of these candidates, in the second phase, we perform full column and cut generation. The size of the branch-and-bound tree is estimated for them according to the approach suggested in (Kullmann, 2009). Then the branching candidate with the smallest estimated tree size is chosen.

4 The case with one available bin per time period and non-negative lags

If we consider the case in which only one bin can be used in each time period ($L = 1$), then the BCP algorithm proposed in Section 3 cannot be used. In this case, given a partition of the item set in bins, the problem of determining if those bins can result in a feasible solution is NP-complete, as shown by Finta and Liu (1996). For this reason, we limit ourselves to the case in which all lags are non-negative numbers. We denote this problem by BPPTL^1_+.

In this context, we can assume that the graph is acyclic. Otherwise it either has a cycle of positive length and then it is trivially infeasible, or it has a cycle with all the arcs having lag zero. In the latter case we can contract all the items in the cycle into one item and have an equivalent problem because these items must be assigned to the same time period, and thus to the same bin.

A problem related to the BPPTL^1_+ has been studied by Kramer et al. (2017), who consider the makespan objective function, i.e., minimizing the last time period in which any bin is used. If all time lags are either zero or one, it is equivalent to minimizing the makespan and the number of bins used. This case generalizes both the bin packing problem with precedence constraints, and the simple assembly line balancing problem, as shown in (Kramer et al., 2017). However, if some time lags are strictly larger than one, minimizing the makespan is not equivalent to the BPPTL^1_+, since a solution that is optimal for one objective function may be sub-optimal for the other objective function. An example of this can be seen in Figure 5. In this example, if we optimize the use of one of the two objective functions we get a solution that is sub-optimal for the other objective function. Minimizing the makespan gives a solution with makespan 5 that uses 4 bins, but minimizing for the number of bins used gives a solution that uses 3 bins with makespan 6.

In this section, we show how to modify the BCP algorithm developed in Section 3 to solve

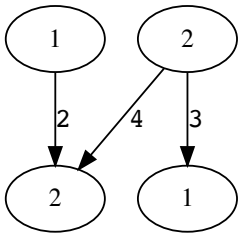


Figure 5: Example of instance with different optimal solutions for different objective functions. Here $W = 2$.

to optimality the BPPTL$^1_+$.

4.1 Valid assignments

Any feasible solution for an instance of the BPPTL$^1_+$ is feasible for the BPPTL$^\infty$, and then the assignment of items to bins (which are also time periods in the BPPTL$^1_+$) must be composed of sets in $\mathcal{B}$, and must also satisfy constraints (7c). However, satisfying these constraints is not sufficient to induce a feasible solution for BPPTL$^1_+$, as can be seen in the example in Figure 6. In this example the time lags force each node to be assigned to the same time period, but the bin capacity does not allow one to assign all items to the same bin. Thus, there is a feasible solution to this instance of the BPPTL$^\infty$, but there are no feasible solutions for the corresponding instance of the BPPTL$^1_+$.

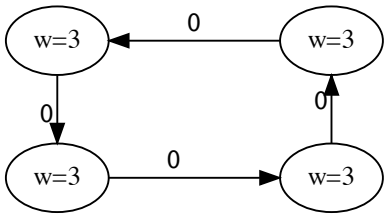


Figure 6: Example of instance with feasible solutions for $BPPTL^\infty$ but not for $BPPTL^1_+$. Capacity $W = 10$

In the following proposition, we characterize feasible solutions for the BPPTL$^1_+$ with a simple condition.

Proposition 4. *A partition $\mathcal{P} \subseteq \mathcal{B}$ induces a feasible solution to the $BPPTL^1_+$ if and only if the aggregated graph $G^{\mathcal{P}}$ is acyclic.*

Proof. Recall that any item set in $\mathcal{P}$ satisfies the capacity constraint. Therefore, only lag constraints have to be checked. If there is a feasible solution, then each set $B \in \mathcal{P}$ can be assigned to a time period p, and then the aggregated graph must be acyclic because otherwise the arc in the cycle with a tail assigned to the largest time period would be violated since all lags are

non-negative. In the other direction, if $\mathcal{G}^{\mathcal{P}}$ is acyclic there is a topological order of the aggregated nodes. A feasible solution can be obtained by assigning bins to time periods following this topological order. □

In the proof, time lag values are not relevant. Then, any partition that induces a feasible solution also induces a feasible solution for any other graph with the same set of arcs, even with different (non-negative) lags. This implies that the time lag values are relevant just to define the conflicts for set $\mathcal{B}$. Thus it only matters if the value of each time lag is zero or a positive number. Therefore, all positive time lags can be replaced by unitary time lags without changing the set of optimal solutions.

In Section 3.1, we define set $\mathcal{C}_R$ to characterize partitions that induce an aggregated graph with a cycle of positive length, and thus an infeasible solution to the BPPTL$^{\infty}$. For the BPPTL$^{1}_{+}$, partitions that induce an aggregated graph with a zero length cycle are also infeasible. Analogously to Definition 2 for cycles of positive length, we now define the set of tuples which induce cycles of length 0 in the aggregated graph.

Definition 3. *For each $R \in \mathbb{N}$, $R \geq 2$, let $\mathcal{C}^0_R \subseteq V^{2R}$ be the set of all tuples $(h_1, h_2, \ldots, h_R, t_1, \ldots, t_R)$ that satisfy the following properties.*

- *All the elements $h_1, \ldots, h_R, t_1, \ldots t_R$ are different.*
- *For each $r \in \{1, \ldots, R\}$, $d(h_r, t_r) \leq 0$ and $d(t_r, h_r) \leq 0$.*
- $\sum_{r=1}^{R} d(t_r, h_{r+1}) = 0$, *where $h_{R+1} = h_1$.*

Recall the definition of value $\theta(B, C)$:

$$\theta(B, C) = \left|\left\{r \in \{1, \ldots, R\} \,:\, \{h_r, t_r\} \subseteq B\right\}\right|.$$

With this definition, constraints (7c) can be rewritten as:

$$\sum_{B \in \mathcal{B}} \theta(B, C)\lambda_B \leq R - 1 \qquad \forall R \geq 2, \forall C \in \mathcal{C}_R. \tag{10}$$

An intuitive extension of these constraints is the following.

$$\sum_{B \in \mathcal{B}} \theta(B, C)\lambda_B \leq R - 1 \qquad \forall R \geq 2, \forall C \in \mathcal{C}^0_R, \tag{11}$$

However, constraints (11) forbid valid solutions. Indeed, if all pairs (h_r, t_r) in a given $C \in \mathcal{C}^0_R$ belong to the same element B in a partition $\mathcal{P}$, this partition does not induce a cycle of length 0 in the aggregated graph.

To overcome this issue, we define for $C \in V^{2R}$ and $B \in \mathcal{B}$ the value $\phi(B, C) = \min\{1, \theta(B, C)\}$. That is, $\phi(B, C)$ is equal to one if any pair in $\mathcal{F}_C$ is contained in B, and is equal to zero otherwise. This definition allows us to characterize partitions that induce an aggregated graph with cycles of length 0.

Proposition 5. *Given a partition $\mathcal{P} \subseteq \mathcal{B}$, the aggregated graph $G^{\mathcal{P}}$ does not have a cycle of zero length if and only if solution λ induced by $\mathcal{P}$ satisfies the following set of constraints:*

$$\sum_{B \in \mathcal{P}} \phi(B, C)\lambda_B \leq R - 1 \qquad \forall R \geq 2, \forall C \in \mathcal{C}^0_R. \tag{12}$$

Proof. Assume that there is a cycle of zero length in $G^{\mathcal{P}}$, and take one with the minimum number of nodes, $B_1, \ldots, B_R$. We take for each $r \in \{1, \ldots, R\}$ nodes $h_r, t_r \in B_r$ such that

$d(t_r, h_{r+1}) = 0$, where $h_{R+1} = h_1$. Then, for $C = (h_1, \ldots, t_R)$ we have $C \in \mathcal{C}_R^0$ due to the minimality of R. For each $r \in \{1, \ldots, R\}$, we have $\phi(B_r, C) = 1$. Then constraint (12) for this tuple C is violated.

In the other direction, let $C \in \mathcal{C}_R^0$ be a tuple such that the respective constraint in (12) is violated. For each $r \in \{1, \ldots, R\}$ there is at most one set $B \in \mathcal{B}$ such that $\lambda_B = 1$ and $\{h_r, t_r\} \subseteq B$. Then there are R different sets B_r, $r = \{1, \ldots, R\}$, such that $\phi(B_r, C) = 1$ and $\lambda_{B_r} = 1$. By definition of $\mathcal{C}_R^0$, there is a cycle of length 0 which contains nodes $G^{\mathcal{P}}$. □

By Proposition 5, the following IP formulation is valid for the BPPTL$_+^1$.

$$\min \sum_{B \in \mathcal{B}} \lambda_B \tag{13a}$$

$$\text{s.t.} \sum_{B \in \mathcal{B}} \mathbb{1}_B(i) \lambda_B = 1 \qquad \forall i \in V, \tag{13b}$$

$$\sum_{B \in \mathcal{B}} \theta(B, C) \lambda_B \leq R - 1 \qquad \forall R \geq 2, \forall C \in \mathcal{C}_R, \tag{13c}$$

$$\sum_{B \in \mathcal{B}} \phi(B, C) \lambda_B \leq R - 1 \qquad \forall R \geq 2, \forall C \in \mathcal{C}_R^0, \tag{13d}$$

$$\lambda_B \in \{0, 1\} \qquad \forall B \in \mathcal{B}. \tag{13e}$$

Here, constraints (13c) forbid cycles of positive length in the aggregated graph, and constraints (13d) forbid cycles of zero length in the aggregated graph.

4.2 Separation

Given a partition $\mathcal{P} \subseteq \mathcal{B}$ induced by an integer solution, a violated constraint can be separated by finding any cycle in the aggregated graph $G^{\mathcal{P}}$. If the length of the cycle is positive, the corresponding constraint (13c) is violated. If the length of the cycle is zero, the corresponding constraint (13d) is violated. The depth-first-search algorithm can be used to find a cycle in the graph $G^{\mathcal{P}}$ or determine that the graph has no cycle. This can be done in time $\mathcal{O}(|\mathcal{P}| + |A^{\mathcal{P}}|)$.

To separate constraints (13d), we adapt the partial enumeration approach developed in Section 3.2.1. The following proposition is useful for this purpose.

Proposition 6. *Let $\left(\bar{\lambda}_B\right)_{B \in \mathcal{B}}$ be a solution satisfying* (13b)*. If multigraph $\mathcal{G} = (V, \mathcal{A})$ constructed from $\bar{\lambda}$ has a cycle $C = (a_1, a_2, \ldots, a_K)$ such that $\sum_{k=1}^K l(a_k) = 0$ and $\sum_{k=1}^K v(a_k) < 1$, and for every $B \in \mathcal{B}$ with $\bar{\lambda}_B > 0$ there is at most one pair $\{i_B, j_B\} \in \mathcal{F}_C$ with $\{i_B, j_B\} \subseteq B$, then there is a constraint* (13d) *violated by $\bar{\lambda}$.*

Proof. Similarly to Proposition 3, we can prove that there is a cycle $C' \subseteq C$ such that constraint $\sum_{B \in \mathcal{B}} \theta(B, C') \lambda_B \leq R - 1$ is violated by $\bar{\lambda}$. For every $B \in \mathcal{B}$ with $\bar{\lambda}_B > 0$ we have $\phi(B, C') = \theta(B, C')$, since at most one of the pairs in $\mathcal{F}_{C'}$ is contained in B. Therefore, constraint (13d) for C' violated by $\bar{\lambda}$. □

The algorithm to search for a violated constraint (13c) or (13d) is similar to Algorithm 1, with two differences. First, when looping over the arcs which leave current node j of current partial path C, we skip arcs $a = (i, j) \in \mathcal{A}_\lambda$ such that there is $B \in \mathcal{B}$ with $\bar{\lambda}_B > 0$ and both $\{i, j\} \subseteq B$ and $\{i', j'\} \subseteq B$ for some $a' = (i', j') \in C \cap \mathcal{A}_\lambda$. Second, in addition to storing cycles with a positive sum of time lags, we also store cycles with the sum of time lags equal to zero.

4.3 Pricing problem

Consider an optimal dual solution $(\bar{\pi}, \bar{\mu}, \bar{\mu}^0)$, with $(\pi_i)_{i \in V}$ of the restricted master problem, where $(\pi_i)_{i \in V}$ are the dual variables of constraints (13b), $(\mu_C)_{C \in \mathcal{C}}$ are the dual variables of

constraints (13c), and $(\mu^0_C)_{C\in\mathcal{C}^0}$ are the dual variables of constraints (13d). The binary coefficient of λ_B in the constraint (13d) corresponding to tuple $C \in \mathcal{C}^0$ is equal to $\phi(B, C)$. Recall that $\phi(B, C)$ is equal to one if and only if there is an item pair $\{i, j\} \in \mathcal{F}_C$ such that both i and j are contained in B. Then the reduced cost of variable λ_B is $1 - \sum_{i\in B} \bar{\pi}_i - \sum_{C\in\mathcal{C}} \theta(B, C)\bar{\mu}_C - \sum_{C\in\mathcal{C}^0} \phi(B, C)\bar{\mu}^0_C$.

We now describe a MIP to find a set B of items with the minimum reduced cost. Let $\bar{\mathcal{C}}$ and $\bar{\mathcal{C}^0}$ be the sets of active constraints (13c) and (13d). Variables x and y are defined in the same ways as for formulation (9). Additionally, a binary variable z_C for each cycle $C \in \bar{\mathcal{C}^0}$ determines whether any of the pairs in $\mathcal{F}_C$ is contained in B. Then the pricing problem can be formulated as the following MIP.

$$\max \sum_{i\in V} \bar{\pi}_i x_i + \sum_{\{i,j\}\in\bar{\mathcal{F}}} \left(\sum_{\substack{C\in\bar{\mathcal{C}}:\\ \{i,j\}\in\mathcal{F}_C}} \bar{\mu}_C \right) y_{i,j} + \sum_{C\in\bar{\mathcal{C}^0}} \bar{\mu}^0_C z_C - 1 \tag{14a}$$

$$\text{s.t.} \sum_{i\in V} w_i x_i \le W, \tag{14b}$$

$$x_i + x_j \le 1 \qquad \forall\{i, j\} \subseteq V,\ d(i, j) \ge 1 \text{ or } d(j, i) \ge 1, \tag{14c}$$

$$x_i + x_j \le 1 + y_{i,j} \qquad \forall\{i, j\} \in \bar{\mathcal{F}}, \tag{14d}$$

$$x_i + x_j \le 1 + z_C \qquad \forall C \in \bar{\mathcal{C}^0}, \forall\{i, j\} \in \mathcal{F}_C, \tag{14e}$$

$$x_i \in \{0, 1\} \qquad \forall i \in V, \tag{14f}$$

$$y_{i,j} \ge 0 \qquad \forall\{i, j\} \in \bar{\mathcal{F}}, \tag{14g}$$

$$z_C \ge 0 \qquad \forall C \in \bar{\mathcal{C}^0}. \tag{14h}$$

Given an optimal solution $(\bar{x}, \bar{y}, \bar{z})$ to (14), we add to the restricted master problem variable $\lambda_{\bar{B}}$, where $\bar{B} = \{i \in V : \bar{x}_i = 1\}$. Since $\mu^0_C < 0$ for $C \in \bar{\mathcal{C}^0}$ and we are maximizing, each variable $\bar{z}_C$ takes the minimum possible value. Therefore, $\bar{z}_C$ is equal to one if and only if $\{i, j\} \subseteq B$ for at least one pair $\{i, j\} \in \mathcal{F}_C$. This implies that z_C is equal to $\phi(B, C)$, and the objective value of solution $(\bar{x}, \bar{y}, \bar{z})$ is equal to the reduced cost of $\lambda_{\bar{B}}$ with the opposite sign.

The proposed formulation for the BPPTL$^1_+$ can be strengthened by reducing collection $\mathcal{B}$ of possible sets of items. Given a triple of items (i, k, j) such that $d(i, k) = 0$ and $d(k, j) = 0$, it is known (Peeters and Degraeve, 2006) that sets B such that $i, j \in B$ and $k \notin B$ can be removed from $\mathcal{B}$. Thus, the following constraints can be added to the pricing problem (14):

$$x_i + x_j \le 1 + x_k, \quad \forall\{i, j, k\} \subseteq V,\ d(i, k) = 0 \text{ and } d(k, j) = 0. \tag{15}$$

Our branch-and-bound algorithm for the KPHSC presented in Section 3.2.3 can be adapted to solve the formulation (14)-(15). For every $C \in \bar{\mathcal{C}^0}$, the value of the lower (primal) bound is adjusted by $\bar{\mu}^0_C$ if at least one of pairs $\{i, j\} \in \mathcal{F}_C$ is contained in V^+. Otherwise, penalty $\bar{\mu}^0_C$ is disregarded in the calculation of the upper (dual) bound as it can only reduce the solution value. Constraints (15) are verified when adding items to set V^-: if $i, j \in V^+$, $d(i, k) = 0$, and $d(k, j) = 0$ then the branch in which item k is added to V^- is not created. Moreover, the best known solution is not updated by the solution formed by items in V^+ if $i, j \in V^+$, $d(i, k) = 0$, $d(k, j) = 0$, and $k \notin V^+$.

5 Computational experiments

We implemented the proposed BCP algorithms in C++ language using a generic branch-cut-and-price library BaPCod (Sadykov and Vanderbeck, 2021). BapCod uses Cplex 12.8 for solving master and pricing subproblems. We also use Cplex 12.8 to solve compact formulation (2) described in Section 2. The experiments were run on a 2 Deca-core Ivy-Bridge Haswell Intel

Xeon E5-2680 v3 server running at 2.50 GHz with 128 GB of RAM. Each instance is solved using a single thread.

5.1 Variant with an unlimited number of available bins

In this section we benchmark different approaches proposed in the paper for the variant of the problem with an unlimited number of available bins. In Section 5.1.1 we describe the test instances which are inspired from the application for the planning of phytosanitary treatments. In Section 5.1.2 we describe a procedure to generate random instances which are based on academic instances of the standard bin-packing problem. All generated instances are freely available at the address `math.u-bordeaux.fr/~rsadykov/#instances`. In Section 5.1.3 we test different cut separation algorithms. Computational comparison of the best variant of the compact formulation and the BCP algorithm is performed in Section 5.1.4. In Section 5.1.5, we compare our BCP algorithm with the standard branch-and-price algorithm based on formulation (3).

5.1.1 Generation of application instances.

The first set of instances we generated are inspired from the application of phytosanitary treatment planning in a vineyard, described in Section 1. Consider a set $\mathcal{Q} = \{1, \dots, Q\}$ of periodic meta-requests for treatment and a planning time horizon $\mathcal{T} = \{1, \dots, T\}$. For each request $q \in \mathcal{Q}$, the ideal elapsed time between two consecutive treatments is denoted by e_q. For additional flexibility, the minimum and maximum elapsed times between two consecutive treatments of request $q \in \mathcal{Q}$ are set to $e_q^- \leq e_q$ and to $e_q^+ \geq e_q$. Assuming that the planning is embedded in a rolling horizon approach, the last treatment before the planning time horizon and the first treatment after the planning time horizon are fixed. Let $f_q \in \{-e_q+1, \dots, 0\}$ be the time period of the last treatment of request $q \in \mathcal{Q}$ before the planning time horizon. Let n_q be the number of treatments of request $q \in \mathcal{Q}$ during the time horizon: it is equal to the maximum integer such that $f_q + n_q e_q \leq T$. Then the time period of the first treatment of request $q \in \mathcal{Q}$ after the planning time horizon is equal to $f'_q = f_q + (n_q + 1) \cdot e_q$. The duration of each treatment i_j^q of request $q \in \mathcal{Q}$ is equal to $w_{i_j^q}$. Treatments are performed by an unlimited fleet of identical vehicles. During a time period, each vehicle is capable of performing treatments of a total duration not exceeding W. The usage cost of a vehicle during one time period is unitary. The objective is to perform all necessary treatments during the planning time horizon while respecting total durations of vehicles and the minimum and maximum elapsed times between any two consecutive treatments of the same request, and to minimize the total vehicle cost.

This problem can be modeled as the BPPTL as follows. Define the bin capacity equal to W and define $L = \infty$. Define two artificial items i^s and i^f with weights $w_{i^s} = w_{i^f} = W$. For each request $q \in \mathcal{Q}$ define n_q items $i_1^q, i_2^q, \dots, i_{n_q}^q$ with weights equal to $w_{i_1^q}, w_{i_2^q}, \dots, w_{i_{n_q}^q}$ respectively. In the time-lag graph, we define arc (i^s, i^f) with lag $T+1$, and arc (i^f, i^s) with lag $-(T+1)$. For each request $q \in \mathcal{Q}$ and each treatment $k \in \{1, \dots, n_q - 1\}$ we define arc (i_k^q, i_{k+1}^q) with lag e_q^-, and arc (i_{k+1}^q, i_k^q) with lag $-e_q^+$. For each $q \in \mathcal{Q}$ we define arc (i^s, i_1^q) with lag $\max\{1, e_q^- + f_q\}$, arc (i_1^q, i^s) with lag $-(e_q^+ + f_q)$, arc $(i_{n_q}^q, i^f)$ with lag $\max\{1, e_a^- - (f'_q - T - 1)\}$, and arc $(i^f, i_{n_q}^q)$ with lag $-(e_a^+ - (f'_q - T - 1))$.

The items created for each request $q \in \mathcal{Q}$ form a double chain in the graph. Each node is connected to the next node in the chain by two arcs, one in each direction. The lags impose the minimum and maximum elapsed time between two consecutive treatments of a request. Items i^s and i^f must be scheduled exactly $T + 1$ time periods apart. Then all other items are scheduled inside T time periods strictly between items i^s and i^f. Thus all treatments are performed within the planning time horizon.

Instances are randomly generated using four integer parameters: the number Q of requests, the number T of time periods, the average number N of request treatments and U the average number of treatments one vehicle can perform in a day. Given tuple (Q, T, N, U), an instance is generated as follows. For each request $q \in \mathcal{Q}$, let e_q be a random integer in the interval

$\left[\left\lfloor 0.7 \cdot \frac{T}{N} \right\rfloor, \left\lceil 1.3 \cdot \frac{T}{N} \right\rceil\right]$, and let $e_q^- = \lfloor 0.8 \cdot e_q \rfloor$ and $e_q^+ = \lceil 1.2 \cdot e_q \rceil$. For each request $q \in \mathcal{Q}$, also let f_q be a random integer in interval $[-e_q + 1, 0]$. The bin capacity W is set to 60, and we define the weight $w_{i_j^q}$ for each treatment of request $q \in A$ as a random integer in interval $\left[\left\lfloor 0.4 \cdot \frac{60}{U} \right\rfloor, \left\lceil 1.6 \cdot \frac{60}{U} \right\rceil\right]$. The dataset we generated consists of 252 instances, one instance for each combination of parameters $Q \in \{3, 5, 7, 9\}$, $T \in \{20, 40, 60, 80, 100, 120\}$, $N \in \{4, 7, 10\}$, and $U \in \{2, 3, 4\}$.

In Figure 1 shown in Section 1, we present the time-lag graph for an instance generated with parameters $Q = 7$, $T = 60$, $N = 4$. Parameter U does not have an impact on the time-lag graph. In this instance, random values generated for the ideal elapsed time of a request are $e_1 = 13$, $e_2 = 11$, $e_3 = 20$, $e_4 = 13$, $e_5 = 14$, $e_6 = 20$, and $e_7 = 14$. Random values generated for the last treatment of a request before the planning time horizon are $f_1 = 0$, $f_2 = -8$, $f_3 = -19$, $f_4 = -12$, $f_5 = -5$, $f_6 = -7$, and $f_7 = -12$.

5.1.2 Generation of random instances.

The instances in the second set we generate are derived from classic bin packing instances by Falkenauer (1996). This set contains eight classes of instances : four classes of "uniform" instances with 120, 250, 500, and 1000 items, and four classes of "triplet" instances with 60, 120, 149, and 501 items. We randomly generate time lags between items according to the following procedure. In order to ensure that all instances are feasible, we randomly generate tentative time period $f_i \in [1, 100]$ for every item $i \in V$. We consider "freedom" parameter F and "density" parameter D. We randomly generate an undirected graph $G' = (V, E)$: edge (i, j) such that $i, j \in V$ belongs to E with probability D. For every edge $(i, j) \in E$, we add lags $l_{i,j} = f_j - f_i - F$ and $l_{j,i} = f_i - f_j - F$. To ensure that all items are scheduled within the time horizon of 100 time periods, we add two artificial items: "source" item 0 and "sink" item $n+1$. Finally we add time lags $l_{0,n+1} = 101$, $l_{n+1,0} = -101$, $l_{0,i} = 1$ and $l_{i,n+1} = 1$ for all $i \in V$.

We take five BPP instances for each class, and for each of these instances, we generate randomly three BPPTL instances which correspond to parameters $(F, D) \in \{(1, 30\%), (3, 10\%), (10, 3\%)\}$. Thus, we generate 15 instances for each of the eight classes, or 120 instances in total.

5.1.3 Comparison of separation approaches.

In Section 3.2.1 two approaches to separate fractional solutions in the BCP algorithm are proposed: solving a MIP and a partial enumeration algorithm. In this section we compare these approaches, as well as the third alternative, which consists in only separating integer solutions. The results for the first set of instances are given in Table 1. Here, column "Gap" presents the average optimality gap, "Root Gap" gives the average gap at the root node of the branch-and-bound tree, "# Opt" the number of instances solved to optimality, "# Nodes" the average number of explored nodes in the branch-and-bound tree, "% Mast" is the time spent to solve the restricted master problem, "% Pric" is the time spent to solve the pricing problem, and "% Cuts" is the time spent to separate and add cuts. The values in the last three columns are in percentage of the total time. The sum of these three values is less than 100%, as the remaining time is due to auxiliary components of BaPCod. The results in Table 1 show that the separation of both integer and fractional solutions by enumeration is the approach that allows us to solve to optimality the largest number of instances, and also to get the smallest average optimality gap over all instances. We can also underline that separating fractional solutions is very important for the efficiency of our BCP algorithm. The pricing problem solution is fast and does not constitute a bottleneck of the BCP algorithm. The majority of time is spent for cut separation, but this time clearly pays off. If cut separation is only used for integer solutions, the number of nodes becomes very large. In this case BaPCod generates a significant overhead, spending more that 50% of the time to manage the search tree. The time spent in the strong diving heuristic is around 30% in general. This time overlaps with the time to solve the restricted master problem and the pricing problem, as well as with the cut separation time, as all these components are used by the heuristic.

Separation approach	Gap	Root Gap	# Opt	# Nodes	% Mast	% Pric	% Cuts
Only integer solutions	4.1%	4.6%	159	13359.6	36.8%	3.5%	0.0%
Fractional solutions by MIP	3.0%	4.4%	180	660.4	27.4%	4.5%	60.9%
Fractional solutions by enum.	2.8%	4.2%	181	420.8	34.5%	7.1%	53.9%

Table 1: Comparison of different separation methods.

5.1.4 Comparison of the BCP algorithm with the Cplex solver.

In this section, we computationally compare two proposed approaches to solve the first set of generated instances with an unlimited number of available bins. The first approach is the best configuration for the compact formulation, which is solved by Cplex (see Appendix B in the online supplement for the results obtained by the different variants, https://pubsonline.informs.org/doi/suppl/10.1287/ijoc.2022. The second approach is our BCP algorithm with the enumeration separation algorithm applied to fractional solutions due to the results we obtained in Section 5.1.3. Both approaches are run with the time limit of one hour.

Table 2 shows the results of this computational comparison. We give results for different subsets of instances grouped by the generation parameters used, as well as the overall results. In the table, column "# Inst." shows the number of instances in the respective subset. Columns under "Opt" show the percentage of the instances that were solved to optimality within the time limit. Columns under "Gap" show the average optimality gap. Column "Root Gap" shows the average optimality gap at the root node. This gap is computed as $\frac{UB-LB}{UB}$, where LB in this case is the lower bound obtained in the root node, UB is the value of the best known solution. Note that "Gap" and "Root Gap" are computed differently, as UB value for "Gap" is the best solution value obtained during the run and not the best known solution value. Thus "Root Gap" can be smaller than "Gap". Columns under "Time" show the average time in seconds. Finally, column "# Nodes" shows the average number of nodes explored in the branch-and-bound tree of the BCP algorithm.

The results, presented in Table 2 show that our BCP algorithm is substantially more efficient than the Cplex solver applied to the compact formulation of the problem. When considering the results for different subsets of instances, we note that the BCP algorithm is not very "sensitive" to the length of the time horizon, whereas the efficiency of Cplex decreases with the increase of parameter T. This is not surprising, as the size of the compact formulation depends heavily on T, whereas the theoretical and practical computational complexity of the components of the BCP algorithm never depend on T. On the other hand, when parameters Q or N increase, both the Cplex and the BCP algorithm become less efficient. The impact of parameter U is however very different. Smaller values of U are better for the BCP algorithm, and larger values are better for Cplex. This behaviour is similar to that observed when solving the classical BPP. When the average number of items which can fit into one bin increases, the quality of lower bounds obtained by the linear relaxation of compact formulations converge to the quality of the column generation lower bound. As the quality of lower bounds become similar, MIP solvers start to be more efficient than column generation approaches due to the fact the linear relaxation of compact formulations can be solved much faster.

5.1.5 Comparison of the BCP algorithm with the standard branch-and-price algorithm.

In the standard branch-and-price algorithm based on formulation (3), we solve the binary knapsack pricing problem with the algorithm by Pisinger (1997). The branching strategy by Vanderbeck (2011) is used. We employ the same stabilization technique and the same diving heuristics as for our BCP algorithm. In Table 3, we show the results of the computational comparison on the second set of instances. The columns are the same as in Table 2. Here, "Gap" and "Root Gap" are calculated only for instances for which both a lower bound and a feasible solution are

	# Inst.	Opt		Gap		Root Gap		Time		# Nodes
		BCP	Cplex	BCP	Cplex	BCP	Cplex	BCP	Cplex	BCP
$T = 20$	36	80.6%	63.9%	1.0%	2.6%	0.9%	8.9%	867.3	1582.3	136.6
$T = 40$	36	75.0%	50.0%	2.0%	6.9%	2.4%	11.6%	920.9	2022.7	253.2
$T = 60$	36	72.2%	44.4%	3.1%	8.5%	3.9%	11.5%	1059.5	2304.0	401.6
$T = 80$	36	69.4%	30.6%	2.9%	10.0%	4.7%	12.7%	1143.8	2851.9	452.9
$T = 100$	36	63.9%	33.3%	4.1%	12.9%	5.3%	13.4%	1418.9	2853.9	686.1
$T = 120$	36	72.2%	16.7%	2.8%	15.8%	6.1%	13.0%	1167.0	3159.3	494.4
$T = 140$	36	69.4%	16.7%	3.5%	16.8%	6.3%	13.3%	1278.8	3097.5	521.1
$Q = 3$	63	100.0%	65.1%	0.0%	6.5%	1.6%	17.8%	72.7	1623.2	21.9
$Q = 5$	63	79.4%	31.7%	2.9%	9.8%	6.1%	11.1%	925.7	2595.7	542.0
$Q = 7$	63	55.6%	19.0%	3.8%	11.1%	5.3%	11.0%	1718.2	3145.2	583.8
$Q = 9$	63	52.4%	30.2%	4.4%	14.5%	4.0%	8.4%	1772.6	2848.3	535.6
$N = 4$	84	100.0%	73.8%	0.0%	2.3%	3.6%	9.2%	71.3	1298.4	62.2
$N = 7$	84	65.5%	31.0%	3.2%	9.2%	4.2%	11.9%	1383.4	2845.4	595.3
$N = 10$	84	50.0%	4.8%	5.1%	20.0%	4.9%	15.0%	1912.2	3515.4	605.0
$U = 2$	84	92.9%	32.1%	0.2%	9.3%	1.4%	9.3%	345.9	2631.8	276.1
$U = 3$	84	66.7%	39.3%	2.6%	10.4%	3.2%	10.8%	1280.8	2368.7	514.7
$U = 4$	84	56.0%	38.1%	5.6%	11.9%	8.2%	16.1%	1740.1	2658.7	471.6
All set 1	252	71.8%	36.5%	2.8%	10.5%	4.2%	12.1%	1122.3	2553.1	420.8

Table 2: Results of the computational comparison between the best variant of the compact formulation solved by Cplex and the best variant of the BCP algorithm

found. We put "-" if no such instances exist for a certain class.

The second set of instances is more difficult to solve than the first set. One can see in Table 3 that BCP algorithm is much more efficient than the branch-and-price algorithm. The main reason is the quality of column generation lower bounds. Another reason is that the restricted master problem based on formulation (3) is harder to solve due to numerous time lag constraints (3c) in the master. For large instances, the column generation procedure cannot even terminate in one hour. In the last row of Table 3, we also compare the BCP and BP algorithms on the first set of application instances. This comparison also shows a clear superiority of the BCP algorithm. We see the quality of column generation lower bounds as the main bottleneck of the BCP algorithm. To improve its efficiency, one should search for new families of strong "non-robust" cutting planes, i.e., cuts which change the structure of the pricing problem.

We do not give detailed results for the Cplex solver applied to the second set of random instances, as the solver performed significantly worse that the BCP algorithm: only one instance from 120 has been solved to optimality, and the run of an absolute majority of remaining instances could not terminate due to insufficient memory.

5.2 Variant with one available bin per time period

In this section we computationally estimate the efficiency of our BCP algorithm for the variant of the BPPTL with one available bin per time period, i.e., for the BPPTL^1_+. We also compare the BCP algorithm to other approaches available in the literature for special cases of the BPPTL^1_+.

5.2.1 Instances.

We tested the branch-cut-and-price algorithm on the BPP-GP instances considered by Kramer et al. (2017). They created a set of instances for the BPP-GP by extending a benchmark set for

	# Inst.	Opt		Gap		Root Gap		Time		Nodes	
		BCP	BP	BCP	BP	BCP	BP	BCP	BP	BCP	BP
$t60$	15	46.7%	26.7%	6.5%	15.1%	7.8%	14.5%	2010.8	2937.8	1463.3	2669.5
$t120$	15	33.3%	0.0%	9.0%	20.7%	8.6%	10.9%	2595.6	>3600	1380.2	3.0
$t249$	15	0.0%	0.0%	6.3%	-	6.3%	7.1%	>3600	>3600	9.1	1.0
$t501$	15	0.0%	0.0%	5.6%	-	5.6%	-	>3600	>3600	1.0	1.0
$u120$	15	20.0%	0.0%	6.7%	17.5%	6.6%	9.9%	2932.6	>3600	1420.1	3.0
$u250$	15	6.7%	0.0%	5.8%	-	5.9%	6.8%	3431.4	>3600	24.6	1.0
$u500$	15	0.0%	0.0%	6.4%	-	6.4%	-	>3600	>3600	1.0	1.0
$u1000$	15	0.0%	0.0%	-	-	-	-	>3600	>3600	1.0	1.0
All set 2	120	13.3%	3.3%	6.7%	16.8%	6.9%	9.8%	3180.7	3517.1	537.5	335.1
All set 1	252	71.8%	39.7%	2.8%	11.3%	4.2%	9.2%	1122.3	2315.7	420.8	5254.6

Table 3: Results of the computational comparison between the best variant of the compact formulation solved by Cplex and the best variant of the BCP algorithm

the SALBP-1 proposed in (Otto et al., 2013). For each instance of the SALBP-1, they created two instances of the BPP-GP by defining random time lags on the arcs of the precedence graph, one of them with random time lag values in $\{0, 1\}$ and the other with random time lag values in $\{0, 1, 2, 3\}$. We will refer to these two special cases of the BPP-GP as BPP-GP01 and BPP-GP03. The literature instances for the BPP-P and the SALBP-1 were also considered by Kramer et al. (2017), as these two problems are special cases of BPP-GP01.

The authors of (Kramer et al., 2017) kindly provided us with the results of their algorithm for the instances of the SALBP-1, the BPP-P, the BPP-GP01, and the BPP-GP03. As mentioned in the introduction, the algorithm in (Kramer et al., 2017) is designed for the makespan objective function. Thus, the problem considered in (Kramer et al., 2017) is a special case of the BPPTL^1_+ only if all time lag values are in $\{0, 1\}$. Thus, we tested our BCP algorithm only on the instances of the SALBP-1, the BPP-P, and the BPP-GP01. Instances containing 20, 50, and 100 items are used. We skip instances with 1000 items as they are out of reach for our algorithm. For each problem variant and each value n equal to the number of items, there are 525 instances in the test set. The structure of precedence graphs is described in (Otto et al., 2013).

Kramer et al. (2017) propose a heuristic procedure to find feasible solutions and thus upper bounds. They also use known techniques from the literature to compute lower bounds. The best upper and lower bounds reported in (Kramer et al., 2017) are usually very good. Most of the instances are solved to optimality. A summary of known results can be seen in Table 4. This table combines the bounds reported in (Kramer et al., 2017), as well as the bounds we obtained using the code from Pereira (2016) provided by the author. Column "# Not Opt" shows the number of instances for which the best known lower bound is strictly smaller than the best known upper bound. Column "# Opt" gives the number of instances for which the best known bounds are equal, i.e., the number of instances with known optimum solutions. In column "Gap Unsolved", the average optimality gap is shown for the open instances.

5.2.2 Experimental results of the BCP algorithm.

In this section, we test our BCP algorithm on the instances described in Section 5.2.1. We use two variants of the BCP algorithm. First variant BCP-K is initialized with the best known upper bounds obtained in (Kramer et al., 2017). Second variant BCP-∞ does not use any initial upper bounds. We impose the time limit of one hour for each instance. Given an instance, we denote by $LB(\text{ALG})$ and $UB(\text{ALG})$ the best lower and upper bounds obtained by algorithm ALG, which can be either BCP-K, BCP-∞, or KDI, where the latter is the approach proposed in (Kramer et al., 2017). In the BCP algorithm, the pricing problem is solved by the custom branch-and-bound algorithm when treating the BPP-GP01 and BPP-P instances. For the SALBP-1

# Items	type	# Not Opt	# Opt	Gap Unsolved
$n = 20$	BPP-GP01	0	525	-
	BPP-P	0	525	-
	SALBP-1	0	525	-
$n = 50$	BPP-GP01	15	510	3.9%
	BPP-P	0	525	-
	SALBP-1	0	525	-
$n = 100$	BPP-GP01	67	458	2.8%
	BPP-P	12	513	2.0%
	SALBP-1	9	516	2.2%

Table 4: Number of instances solved to optimality in (Kramer et al., 2017) and average optimality gap for unsolved instances

instances, the MIP solver is used to solve the pricing problem. This choice is explained by the fact that some SALBP-1 instances have a large bin capacity and very few hard conflicts, and our branch-and-bound algorithm is not efficient for such instances of the knapsack problem with hard and soft conflicts.

In Table 5 we compare upper bounds UB(BCP-∞) and UB(KDI). Instances are divided into three groups: "Better", "Equal" and "Worse", depending on whether UB(KDI)$>$ UB(BCP-∞), UB(KDI)$=$ UB(BCP-∞), and UB(KDI)$<$ UB(BCP-∞). Columns under "Avg. Diff." give the average absolute difference between values UB(KDI) and UB(BCP-∞) for the corresponding group of instances. For some groups of instances, this average difference is marked with "∞", which means that for at least one instance in this group no feasible solution was found by algorithm BCP-∞.

		Better		Equal	Worse	
# Items	type	#	Avg. Diff.	#	#	Avg. Diff.
$n = 20$	BPP-GP01	0	-	525	0	-
	BPP-P	0	-	525	0	-
	SALBP-1	0	-	525	0	-
$n = 50$	BPP-GP01	1	1.0	520	4	1.0
	BPP-P	0	-	501	24	1.0
	SALBP-1	0	-	515	32	1.0
$n = 100$	BPP-GP01	13	1.0	480	32	∞
	BPP-P	2	1.0	372	151	1.5
	SALBP-1	0	-	339	186	∞

Table 5: Comparison of upper bounds UB(BCP-∞) and UB(KDI)

The results in Table 5 show that the heuristic proposed in (Kramer et al., 2017) is usually much better to obtain good feasible solutions than the algorithm BCP-∞. Nevertheless, our algorithm is able to improve the best known solutions for 16 instances. Using the variant BCP-K of our algorithm, no more best known solutions were improved.

In Table 6, we compare lower bounds LB(BCP-K) and LB(KDI). In the same vein as above, the instances are divided into the groups "Better", "Equal" and "Worse", depending on whether LB(KDI)$<$ LB(BCP-K), LB(KDI)$=$ LB(BCP-K), and LB(KDI)$>$ LB(BCP-K). The results show that the lower bounds obtained by our algorithm BCP-K are on average worse than the best known ones for instances of the BPP-P and the SALBP-1. However, our lower bounds are on average better than the best known ones for instances of the BPP-GP01. This is not surprising as the BPP-GP has been less studied in the literature. Overall, our algorithm improved 63 best

known lower bounds among 103 open instances, i.e., for the majority of open instances.

# Items	type	Better #	Better Avg. Diff.	Equal #	Worse #	Worse Avg. Diff.
$n = 20$	BPP-GP01	0	-	525	0	-
	BPP-P	0	-	525	0	-
	SALBP-1	0	-	525	0	-
$n = 50$	BPP-GP01	13	1.0	511	1	1.0
	BPP-P	0	-	501	24	1.0
	SALBP-1	0	-	522	3	1.0
$n = 100$	BPP-GP01	36	1.4	472	17	1.1
	BPP-P	8	1.0	448	69	1.2
	SALBP-1	6	1.0	487	32	1.0

Table 6: Comparison of lower bounds LB(BCP-K) and LB(KDI)

In Table 7, we show the optimality status of instances before and after applying our algorithm. Column "# Opt KDI" shows the number of instances with known optimum solutions in the literature. Column "# Opt. +BCP" gives the number of instances with known optimality based on the literature results and the results obtained by our BCP algorithm. Finally, "# Not Opt." shows the number of instances that still remain open. This table shows that for 79 instances the optimality status is obtained for the first time among 103 open instances. 24 instances still remain open.

# Items	type	# Opt KDI	# Opt. +BCP	# Not Opt.
$n = 20$	BPP-GP01	525	0	0
	BPP-P	525	0	0
	SALBP-1	525	0	0
$n = 50$	BPP-GP01	510	14	1
	BPP-P	525	0	0
	SALBP-1	525	0	0
$n = 100$	BPP-GP01	458	49	18
	BPP-P	513	10	2
	SALBP-1	516	6	3

Table 7: Optimality status of instances

Finally, in Table 8, we give statistics for our algorithm BCP-K applied to instances described in Section 5.2.1. Columns under "# Nodes" show the average number of explored nodes in the branch-and-bound tree for both unsolved instances and instances solved to optimality. Column "Not Solved" gives the number of instances not solved to optimality by the algorithm BCP-K. Column "Solved" shows the number of instances solved to optimality. Column "Time" gives the average time in seconds our algorithm took when it could solve instances to optimality. It can be seen from the results in Table 8 that algorithm BCP-K is less efficient than the approach by Kramer et al. (2017) for the BPP-P and the SALBP-1, but more efficient for the BPP-GP01. We should mention, however, that our algorithm is initialized by the best known solutions obtained in (Kramer et al., 2017). Thus, it makes sense to combine our BCP algorithm with the approach proposed in (Kramer et al., 2017) to solve instances of the BPP-GP01. These two methods seem to have a complementary strength, according to the results presented in Table 7.

# Items	type	# Nodes		# Instances		Time
		Not Solved	Solved	Not Solved	Solved	Solved
$n = 20$	BPP-GP01	-	1.5	0	525	0.19
	BPP-P	-	1.7	0	525	0.18
	SALBP-1	-	1.2	0	525	0.84
$n = 50$	BPP-GP01	8049.0	15.6	2	523	6.59
	BPP-P	3479.4	40.3	24	501	36.92
	SALBP-1	621.7	4.5	3	522	25.64
$n = 100$	BPP-GP01	5969.1	26.7	34	491	34.71
	BPP-P	2491.7	36.4	71	454	57.70
	SALBP-1	52.2	2.2	35	490	182.20

Table 8: Statistics of the BCP-K algorithm

6 Conclusions

In this work, we introduce the bin packing problem with time lags (BPPTL). As a motivation, we describe an application of the BPPTL in which one needs to decide on the planning of phytosanitary treatments in a vineyard. We first present an IP formulation for the problem. Then for two important special cases of the BPPTL with an unlimited number of available bins and with one available bin per time period, we propose an exact branch-cut-and-price (BCP) algorithm. This algorithm relies on a known IP formulation with an exponential number of variables which define a relaxation for the problem, and a new family of constraints which serve to cut infeasible solutions produced by the relaxation. We present two approaches to separate the new family of constraints. The BCP algorithm also incorporates an automatic dual price smoothing technique to improve the convergence of column generation, a strong diving heuristic for finding feasible solutions, and a strong multi-phase Ryan&Foster branching.

The computational experiments show that our BCP algorithm substantially outperforms a commercial MIP solver on newly generated instances inspired from the phytosanitary treatment planning application. The experiments on literature instances for special cases of the BPPTL show that the BCP algorithm can complement known approaches, as it allows us to obtain the optimality status for 70% of previously open instances.

Nevertheless, we believe that the efficiency of our BCP algorithm can still be improved. The experiments on literature instances show that the primal heuristic should be improved. The BCP algorithm currently relies on the generic strong diving heuristic, which can be slow and relatively inefficient for larger instances. Therefore, developing specialized heuristics for the BPPTL is an important research direction. Known heuristics for the related problem BPP-GP, such as the one proposed in (Kramer et al., 2017), may serve as a basis for this work.

The BCP algorithm is currently limited to two special cases of the BPPTL. Extending it to the general case is an interesting research perspective. This seems not to be easy to do, as in general it is an NP-complete problem (Finta and Liu, 1996) to determine whether a partition of items into bins is feasible or not with respect to the time lags. Thus, any family of cutting planes to cut off infeasible assignments of items to bins will be NP-hard to separate even if only integer assignments are considered. An extension to the general case is relevant in practice, as sometimes we cannot assume that we may use any number of resources as we want in any time period.

Another important generalization concerns related scheduling problems. Release dates and deadlines can be modeled with time lags, and thus addressed easily. However, considering objective functions like makespan or total tardiness, and therefore time period-dependent penalization of item placement, would require a substantial modification of the algorithm.

In the phytosanitary treatment planning application, we assume that the time needed for a

vehicle to go from one sector to another is negligible in comparison to the duration of treatments. Thus transition times of vehicles are ignored. In practice, this is not always true. If transition times are not negligible, the problem shifts from the class of precedence-constrained bin packing problems to the class of periodic vehicle routing problems (Campbell and Wilson, 2014).

In the standard periodic vehicle routing problem (PVRP), there are only a few alternatives for visiting each client. For example, a client may be visited on Monday and Thursday, or on Tuesday and Friday. If the planning time horizon is sufficiently long, having only a few alternatives dramatically limits the flexibility of visits. Flexible variants of the PVRP exist, such as one proposed in (Archetti et al., 2017). However, flexibility in that case concerns the amount of goods delivered to a client during each visit. An extension of the BPPTL to a flexible periodic vehicle routing problem in which one imposes minimum and maximum time lags on consecutive visits to the same client has a large number of applications in the context of provision of services such as periodic maintenance visits. A BCP algorithm similar to the one proposed in this paper may be efficient for such a flexible periodic vehicle routing problem, provided that the pricing problem takes into account the vehicle transition times.

Finally, the BPPTL structure is sufficiently rich that there may be other applications which result in time lag graphs in a class different from the class of double-linked chain graphs. Finding such applications and suggesting test instances for them is useful for further progress in solution approaches for the BPPTL.

Acknowledgments

We thank Raphael Kramer for kindly providing us with results for individual instances and other details on their work.

We thank Jordi Pereira for kindly providing us with results for individual instances as well as his code for the BPP-P.

Experiments presented in this paper were carried out using the PlaFRIM (Federative Platform for Research in Computer Science and Mathematics), created under the Inria PlaFRIM development action with support from Bordeaux INP, LABRI and IMB and other entities: Conseil Régional d'Aquitaine, Université de Bordeaux, CNRS and ANR in accordance to the "Programme d'Investissements d'Avenir" (see `www.plafrim.fr/en/home`).

This work was supported by the Government of Chile through CONICYT-PFCHA/Doctorado Nacional/2017-211713157 (ORL).

References

Tobias Achterberg. *Constraint Integer Programming.* PhD thesis, Technische Universitat Berlin, 2007.

Claudia Archetti, Elena Fernández, and Diana L. Huerta-Muñoz. The flexible periodic vehicle routing problem. *Computers & Operations Research*, 85:58 – 70, 2017.

Andrea Bettinelli, Valentina Cacchiani, and Enrico Malaguti. A branch-and-bound algorithm for the knapsack problem with conflict graph. *INFORMS Journal on Computing*, 29(3):457–473, 2017.

Andrew Bilyk, Lars Mönch, and Christian Almeder. Scheduling jobs with ready times and precedence constraints on parallel batch machines using metaheuristics. *Computers & Industrial Engineering*, 78:175–185, 2014.

Ann Melissa Campbell and Jill Hardin Wilson. Forty years of periodic vehicle routing. *Networks*, 63(1):2–15, 2014.

Stefano Coniglio, Fabio Furini, and Pablo San Segundo. A new combinatorial branch-and-bound algorithm for the knapsack problem with conflicts. *European Journal of Operational Research*, 289(2):435–455, 2021.

Mauro Dell'Amico, José Carlos Díaz Díaz, and Manuel Iori. The Bin Packing Problem with Precedence Constraints. *Operations Research*, 60(6):1491–1504, 2012. doi: 10.1287/opre.1120.1109. URL https://doi.org/10.1287/opre.1120.1109.

Maxence Delorme and Manuel Iori. Enhanced pseudo-polynomial formulations for bin packing and cutting stock problems. *INFORMS Journal on Computing*, 32(1):101–119, 2020.

Maxence Delorme, Manuel Iori, and Silvano Martello. Bin packing and cutting stock problems: Mathematical models and exact algorithms. *European Journal of Operational Research*, 255 (1):1–20, 2016.

Lionel Dupont and Clarisse Dhaenens-Flipo. Minimizing the makespan on a batch machine with non-identical job sizes: an exact procedure. *Computers & Operations Research*, 29(7):807–819, 2002.

Samir Elhedhli, Lingzi Li, Mariem Gzara, and Joe Naoum-Sawaya. A Branch-and-Price Algorithm for the Bin Packing Problem with Conflicts. *INFORMS Journal on Computing*, 23(3): 404–415, 2011.

Simon Emde, Lukas Polten, and Michel Gendreau. Logic-based benders decomposition for scheduling a batching machine. *Computers & Operations Research*, 113:104777, 2020.

Emanuel Falkenauer. A hybrid grouping genetic algorithm for bin packing. *Journal of Heuristics*, 2:5–30, 1996.

Albert E. Fernandes Muritiba, Manuel Iori, Enrico Malaguti, and Paolo Toth. Algorithms for the Bin Packing Problem with Conflicts. *INFORMS Journal on Computing*, 22:401–415, 2010.

Lucian Finta and Zhen Liu. Single machine scheduling subject to precedence delays. *Discrete Applied Mathematics*, 70(3):247–266, 1996. ISSN 0166-218X. doi: https://doi.org/10.1016/0166-218X(96)00110-2. URL http://www.sciencedirect.com/science/article/pii/0166218X96001102.

Michel Gendreau, Gilbert Laporte, and Frédéric Semet. Heuristics and lower bounds for the bin packing problem with conflicts. *Computers and Operations Research*, 31(3):347 – 358, 2004.

C Joncour, S Michel, R Sadykov, D Sverdlov, and F Vanderbeck. Column Generation based Primal Heuristics. *Electronic Notes in Discrete Mathematics*, 36:695–702, 2010. ISSN 1571-0653. doi: https://doi.org/10.1016/j.endm.2010.05.088. URL http://www.sciencedirect.com/science/article/pii/S1571065310000892.

Leonid V Kantorovich. Mathematical methods of organizing and planning production. *Management science*, 6(4):366–422, 1960. Translation of a 1939 paper in Russian.

Ali Khanafer, François Clautiaux, and El-Ghazali Talbi. New lower bounds for bin packing problems with conflicts. *European Journal of Operational Research*, 206(2):281 – 288, 2010.

Ali Khanafer, François Clautiaux, and El-Ghazali Talbi. Tree-decomposition based heuristics for the two-dimensional bin packing problem with conflicts. *Computers & Operations Research*, 39(1):54 – 63, 2012.

Raphael Kramer, Mauro Dell'Amico, and Manuel Iori. A batching-move iterated local search algorithm for the bin packing problem with generalized precedence constraints. *International Journal of Production Research*, 55(21):6288–6304, 2017. doi: 10.1080/00207543.2017.1341065. URL https://doi.org/10.1080/00207543.2017.1341065.

O Kullmann. *Handbook of Satisfiability*, chapter Fundaments of branching heuristics, pages 205–244. IOS Press, Amsterdam, 2009.

David R Morrison, Edward C Sewell, and Sheldon H Jacobson. An application of the branch, bound, and remember algorithm to a new simple assembly line balancing dataset. *European Journal of Operational Research*, 236(2):403–409, 2014. ISSN 0377-2217. doi: https://doi.org/10.1016/j.ejor.2013.11.033. URL http://www.sciencedirect.com/science/article/pii/S0377221713009508.

Alena Otto, Christian Otto, and Armin Scholl. Systematic data generation and test design for solution algorithms on the example of SALBPGen for assembly line balancing. *European Journal of Operational Research*, 228(1):33–45, 2013. ISSN 0377-2217. doi: https://doi.org/10.1016/j.ejor.2012.12.029. URL http://www.sciencedirect.com/science/article/pii/S0377221713000039.

Marc Peeters and Zeger Degraeve. An linear programming based lower bound for the simple assembly line balancing problem. *European Journal of Operational Research*, 168(3):716–731, 2006.

Jordi Pereira. Empirical evaluation of lower bounding methods for the simple assembly line balancing problem. *International Journal of Production Research*, 53(11):3327–3340, 2015. doi: 10.1080/00207543.2014.980014. URL https://doi.org/10.1080/00207543.2014.980014.

Jordi Pereira. Procedures for the bin packing problem with precedence constraints. *European Journal of Operational Research*, 250(3):794–806, 2016. ISSN 0377-2217. doi: https://doi.org/10.1016/j.ejor.2015.10.048. URL http://www.sciencedirect.com/science/article/pii/S0377221715009741.

Artur Pessoa, Ruslan Sadykov, Eduardo Uchoa, and François Vanderbeck. Automation and combination of linear-programming based stabilization techniques in column generation. *INFORMS Journal on Computing*, 30(2):339–360, 2018.

Artur Pessoa, Ruslan Sadykov, Eduardo Uchoa, and François Vanderbeck. A generic exact solver for vehicle routing and related problems. *Mathematical Programming B*, 183:483–523, 2020.

David Pisinger. A minimal algorithm for the 0-1 knapsack problem. *Operations Research*, 45(5): 758–767, 1997.

N. Rafiee Parsa, B. Karimi, and A. Husseinzadeh Kashan. A branch and price algorithm to minimize makespan on a single batch processing machine with non-identical job sizes. *Computers & Operations Research*, 37(10):1720–1730, 2010.

D. M. Ryan and B. A. Foster. An integer programming approach to scheduling. In A. Wren, editor, *Computer Scheduling of Public Transport Urban Passenger Vehicle and Crew Scheduling*, pages 269 – 280. North-Holland, Amsterdam, 1981.

Ruslan Sadykov and François Vanderbeck. Bin Packing with Conflicts: A Generic Branch-and-Price Algorithm. *INFORMS Journal on Computing*, 25(2):244–255, 2013. doi: 10.1287/ijoc.1120.0499. URL http://dx.doi.org/10.1287/ijoc.1120.0499.

Ruslan Sadykov and François Vanderbeck. BaPCod — a generic Branch-And-Price Code. Technical report HAL-03340548, Inria Bordeaux — Sud-Ouest, September 2021.

Ruslan Sadykov, François Vanderbeck, Artur Pessoa, Issam Tahiri, and Eduardo Uchoa. Primal heuristics for branch-and-price: the assets of diving methods. *INFORMS Journal on Computing*, 31(2):251–267, 2019.

Armin Scholl and Christian Becker. State-of-the-art exact and heuristic solution procedures for simple assembly line balancing. *European Journal of Operational Research*, 168(3):666–693, 2006. ISSN 0377-2217. doi: https://doi.org/10.1016/j.ejor.2004.07.022. URL http://www.sciencedirect.com/science/article/pii/S0377221704004795.

R. Uzsoy. Scheduling a single batch processing machine with non-identical job sizes. *International Journal of Production Research*, 32(7):1615–1635, 1994.

R. Uzsoy. Scheduling batch processing machines with incompatible job families. *International Journal of Production Research*, 33(10):2685–2708, 1995.

François Vanderbeck. Branching in branch-and-price: a generic scheme. *Mathematical Programming*, 130(2):249–294, 2011.

Laguna Wei, Zhixing Luo, Roberto Baldacci, and Andrew Lim. A new branch-and-price-and-cut algorithm for one-dimensional bin-packing problems. *INFORMS Journal on Computing*, 32 (2):428–443, 2020.