



HAL
open science

A dedicated pricing algorithm to solve a large family of nurse scheduling problems with branch-and-price

Antoine Legrain, Jérémy Omer

► To cite this version:

Antoine Legrain, Jérémy Omer. A dedicated pricing algorithm to solve a large family of nurse scheduling problems with branch-and-price. *INFORMS Journal on Computing*, 2024, 10.1287/ijoc.2023.0019 . hal-03964952

HAL Id: hal-03964952

<https://hal.science/hal-03964952v1>

Submitted on 31 Jan 2023

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.

A dedicated pricing algorithm to solve a large family of nurse scheduling problems with branch-and-price

ANTOINE LEGRAIN

Polytechnique Montreal, CIRRELT, GERAD, Montreal, QC, Canada

antoine.legrain@polymtl.ca

JÉRÉMY OMER

Univ Rennes, INSA Rennes, CNRS, IRMAR - UMR 6625, 35000 Rennes, France

jeremy.omer@insa-rennes.fr

Abstract

In this paper, we describe a branch-and-price algorithm for the personalized nurse scheduling problem. The variants that appear in the literature involve a large number of constraints that can be hard or soft, meaning that they can be violated at the price of a penalty. We capture the diversity of the constraints on individual schedules by seven generic constraints characterized by lower and upper bounds on a given quantity. The core of the column generation procedure is in the identification of individual schedules with minimum reduced cost. For this we solve a shortest path problem with resource constraints (SPPRC) where several generic constraints are modeled as resource constraints. We then describe dominance rules adapted to the presence of both upper and lower bounds on the resources and leverage soft constraints to improve the dominance. We also describe several acceleration techniques for the solution of the SPPRC, and branching rules that fit the specificities of the problem. Our numerical experiments are based on the instances of three benchmarks of the literature including those of the two international nurse rostering competitions (INRC-I and INRC-II). Their objective is three-fold: assess the dominance rules and the acceleration techniques, investigate the capacity of the algorithm to find provable optimal solutions of instances that are still open, and conduct a comparison to best published results. The most noticeable conclusion is that the improved solution of the SPPRC allows to solve optimally all the INRC-II instances where a 4-weeks planning horizon is considered and 40% of the 8-weeks instances.

1 Introduction

1.1 The nurse scheduling problem

Personalized nurse scheduling problem (NSP) is more important than ever in this post-pandemic context, where important nurses shortages still exist and have even increased during the pandemic in most western healthcare systems. Proposing high-quality schedules that fit the nurses preferences while ensuring the right coverage is a key-element to help managers retain nurses in their care units. Moreover, these schedules need to cover an horizon big enough, typically 4 to 8 weeks, to let the nurses plan their own personal time based on their work schedules.

The NSP is the design of the schedule of a nursing service so as to satisfy the needs of the nurses and all operational constraints at minimal cost. A day is divided into several time frames called *shifts*, and each nurse has a set of *skills* corresponding to tasks she can perform. When a nurse works on a given day, she or he is assigned to one shift with a skill he or she possesses. A day with no assignment is a *day off*. The NSP then consists in building a *roster* for each nurse, i.e., a sequence of assignments and days off covering the planning horizon. Most rostering problems aim at balancing the work among employees and ensuring that the personal rosters satisfy working regulations and contracts. As a consequence the cost of the complete schedule usually aggregates an evaluation of the global quality of service it provides with penalties associated to unbalanced personalized rosters and to the violations of contracts and working regulations. The quality of service is in general quantified by a demand on each shift of each day, and a possibly large number of constraints define the requirements of each personalized roster.

The review article by [Che+03] lists 16 types of constraints that are commonly found in the literature. For instance, the workload of each nurse usually has to lie between a minimum and a maximum number of assignments on the planning horizon. It is also usual that the number of consecutive assignments be lower and upper bounded. Depending on the considered country and hospital, each constraint may be *hard*, meaning that a feasible solution

necessarily satisfies it, or *soft*, in which case it can be violated at the price of a linear penalty. As an example, the two international nurse rostering competitions (INRC), INRC-I [Has+14] and INRC-II [Ces+19], specify two rostering problems where most constraints are soft. The hard constraints of the INRC instances specify that nurses can perform only one assignment per day, that they need the corresponding skill to perform a task and that a minimum number of nurses must be assigned to each shift. Up to 15 other soft constraints then characterize the required working conditions of each nurse. It should be noted that the INRC-II problem originally considered a stochastic NSP where the demands are dynamically revealed on a weekly basis, but the literature has often focused on its *static* version where all data is known in advance.

A fair deal of the complexity of the NSP can be attributed to the personalized aspect of the schedules: the same schedule may indeed have very different costs if given to one nurse or another, because they have different contracts or preferences on days off. In contrast, workforce scheduling problems usually build anonymous rosters thus making the problem easier.

1.2 Solution of the nurse scheduling problem

Most approaches that were used to tackle NSP instances are based on metaheuristics or mathematical programming techniques. Among them, several methods solely based on metaheuristics have shown their capacity to compute good rosters in a small computational time on a large variety of benchmarks (see *e.g.* [Bur+13; CGS20; Abu+21; Abd+21]). However, the results of the two INRCs emphasized that these approaches were outperformed by methods based on integer programming formulations of the problem [CGS20; Has+14] except for instances with very long horizons (16 and 32 weeks) [CGS20]. Another limit of these approaches is that they don't provide lower bounds on the optimal value, so they cannot prove the optimality of the solutions by themselves.

The mathematical programming approaches can be categorized in two families. In the first one, the problem is formulated as a compact integer linear program (ILP) where the variables refer to working assignments and resting days. This has been investigated by Römer [Röm16] to successfully handle the stochastic problem described in INRC-II. Santos et al. [San+16b] have also been able to find provable optimal solutions of many INRC-I instances and improve on the optimality gap of others. However, most hidden instances remain open.

The method that we describe in this article belongs to the second family, where the problem is modeled as an extended formulation where variables refer to individual rosters or sequences of working assignments. Given that this formulation includes in general an exponential number of variables, its solution relies on a *column generation* procedure embedded in a *branch-and-price* algorithm to get integer solutions. These methods, possibly coupled with matheuristics, have been able to produce the best results on most benchmarks. Below, we draw a detailed review of the previously developed branch-and-price approaches for the NSP and related problems, with a particular focus on the computation of individual rosters.

Branch-and-price approaches

In most extended formulations of the NSP, the columns of the constraint matrix correspond to rosters. The linear constraints are partitioning equalities or covering inequalities, each of them ensuring that a given task is performed by a sufficient number of nurses. Due to the length of the planning horizon and the number of different tasks that may be performed by each nurse, it is impossible to enumerate all columns beforehand and solve the resulting ILP with integer programming algorithms such as branch-and-cut. To palliate this numerical obstacle, one restricts the number of columns in the problem at first and, at each node of the branch-and-bound tree, the linear relaxation is solved by column generation; the resulting algorithm is called branch-and-price. For more details on the implementation of branch-and-price techniques, the reader is referred to [Bar+98]. In scheduling applications, the generation of new columns is usually based on the solution of a shortest path problem with resource constraints (SPPRC). For this, one considers a network where vertices represent the tasks to complete and arcs the possible transitions. Each arc is weighted with a cost and its traversal consumes certain resources. The SPPRC then consists in finding the minimum cost path between two artificial origin and sink vertices such that the resources consumed along the path lie within given bounds. Its solution either yields a new column of negative reduced cost or asserts that the optimal solution of the restricted linear relaxation is optimal for the non-restricted version. The reader can find an overview of column-generation algorithms in [DDS06] and a comprehensive review on SPPRC in [ID+05].

The first published attempt to tackle the NSP by means of branch-and-price dates back from 1998 [JSV98]. The authors develop a generic framework that handles a large variety of constraints. New columns are generated by solving an SPPRC that involves up to seven different constraints. Bard and Purnomo [BP05] stress the personalized facet of the problem. This leads them to generate new schedules by a swap heuristic and check their feasibility a posteriori. In [MV10] a heuristic is first run as an attempt to generate negative reduced cost rosters and the

SPPRC is solved only when the attempt fails. They also compare more than fifteen different branching rules on one or several variables. Burke and Curtois [BC14], generate columns by solving an SPPRC with a heuristic dynamic programming technique. They report excellent results on the INRC-I benchmark. The pricing problem can also be solved with constraint programming, as in [HQ12]. Gomes, Toffolo, and Santos [GTS17] enrich their branch-and-price scheme with a variable neighborhood search to tackle static INRC-II instances with a 4-weeks horizon. They use a classical decomposition by roster for each nurse.

Another decomposition of the problem draws its inspiration from airline crew scheduling problems. For instance, in [KK04], each column is a rotation, *i.e.*, a sequence of flights that starts and ends at the same base. The rostering problem then consists in choosing a sequence of rotations to build a roster for each crew member. Legrain, Omer, and Rosat [LOR20] adapted this idea to the NSP by defining a rotation as a sequence of working assignments preceded and followed by a sequence of days off. They enriched their branch-and-price algorithm with a large neighborhood search to find good feasible solutions in limited time. This approach produced the best results for a large number of INRC-II instances, as shown in the recent experimental comparison carried out by Ceschia, Guido, and Schaefer [CGS20]. In his PhD thesis, Lensing [Len20] reports that his implementation of this rotation decomposition is outperformed by the classical roster decomposition on the benchmark proposed by Curtois and Qu [CQ14], where most constraints are hard.

Generation of individual rosters A key step of the branch-and-price algorithms for NSP is the generation of individual rosters with negative reduced costs. In our previous works on the subject, we have observed that this is the part of the algorithm where most of the computational effort is spent. In most works of the literature, this is done by solving an SPPRC in a graph where the nodes stand for possible assignments or resting days and where the arcs represent allowed transitions. The specificity of the NSP is that resource constraints may be soft and that they specify lower and upper bounds on the same quantity. While describing the first dynamic programming approach to the SPPRC, Joksich [Jok66] mentioned that in the presence of two inequalities imposing upper and lower limits on a resource consumption, a path can be identified as less efficient than another only if the resource consumption is the same along both paths. He then stressed that dynamic programming approaches would certainly not be computationally worthwhile unless many paths share the same resource values. Later, the dynamic programming recursion proposed by Beasley and Christofides [BC89] also kept one path for each possible value of the resources, thus leading to the same computational issue.

The research on SPPRC has often been driven by the study of vehicle routing problems (VRPs). In this context, when both lower and upper bounds are considered, they are related to time windows $[a_i, b_i]$ where each node i must be visited. In this application, it is usually possible to wait at a node before serving it. As a consequence, the lower bound of a time window will not impact the validity of a path arriving at node i earlier than a_i . Instead, it impacts the departure time from node i : a vehicle reaching i before a_i will wait until a_i to serve it and then leave node i . Classical time-windows are enforced as hard constraints, but some works considered soft time windows. In the latter variant, waiting is still possible without extra cost at any node, but if a node is served earlier or later than its time window a linear penalty must be paid. A section is devoted to soft time windows in the review article by Costa, Contardo, and Desaulniers [CCD19]. Sexton and Choi [SC86] were the first to consider soft time windows in a VRP. Those were introduced to distinguish working time windows, outside of which a customer could not be served, and the preference time slots of the customers, which can be violated by paying a penalty. As emphasized by Liberatore, Righini, and Salani [LRS11], the greatest difficulty that they raise is in the price of earliness. A compromise has to be found between earliness penalty and tardiness in subsequent nodes, hence dominance happens a lot less when solving the subproblems with label-setting algorithms. To mitigate this issue, Liberatore, Righini, and Salani associate a non-increasing linear stepwise cost function with each label to represent an infinite number of states where the other resources' consumption and the list of visited nodes are the same. This prevents from having one label for each possible waiting time at each node. A similar approach is followed by [BCR14] and [AJ14].

Unfortunately, given that waiting is not relevant in general for resource constraints, the above cannot be generalized to nurse scheduling. In contrast, Tagmouti, Gendreau, and Potvin [TGP07] have studied a variant of the traveling salesman problem where a time-dependent service cost is given by a convex function of the arrival time at each node. They solve the problem with a column-generation algorithm where the subproblem is a shortest elementary path problem with one resource constraint. Their algorithm is a direct adaptation of the classical label-setting algorithm where the cost of a partial path is given by the sum of the arc traveling costs and the service costs. In [QFY10b], early arrivals are penalized linearly. The authors clearly emphasize that soft constraints result in a dramatic increase of the computational time that is required to find optimal paths, especially when early arrivals are penalized. But in every case, the soft constraints yield no more adaptation than in the computation of the cost of a partial path. When early arrivals are penalized, the shortest path problem is solved with a genetic algorithm

heuristics [QTY10a].

In the context of NSP, only a few works have handled soft constraints in a branch-and-price approach where the pricing problem is solved by a label-setting algorithm. For instance, Strandmark, Qu, and Curtois [SQC20] developed a heuristic column generation approach, but the algorithm does not produce lower bounds because their SPPRC does not capture the soft roster constraints. In contrast, Burke and Curtois [BC14] straightforwardly adapted the classical dominance rule to handle lower bounds: a label can dominate another label only if the consumption of every upper bounded resource is smaller or equal and if the consumption of every lower-bounded resource is larger or equal. As a consequence, dominance happens only when the consumption is equal for the resources that are both upper and lower-bounded. In [LOR20], the network of the SPPRC is modified to reduce the pricing problem to a standard SPPRC with only one upper bounded resource. For this, short sequences of assignments are enumerated, and the nodes and arcs are duplicated into layers where the number of consecutive assignments to the same shift can be tracked. Dominance is much more efficient with this approach, but the network is much larger. Moreover, dominance is not cross-checked between the nodes that belong to different layers but correspond to the same assignment.

1.3 Main contributions

The purpose of the research described in this article is to implement a branch-and-price algorithm that can handle most of the NSP instances described in the literature while finding provable optimal solutions for a large number of instances that are still open. Up to now, most static INRC-II instances studied in the literature were still open, and so were some INRC-I instances. Moreover, most studies have focused on specific instances and either the corresponding code has not been shared publicly or it is dedicated to the specific constraints considered in these instances.

Algorithmic contributions In the development of our algorithm, we have looked into opportunities of improvements in every components of the branch-and-price algorithm, but we have drawn a particular focus on the generation of individual rosters. As a consequence, we have designed a new dominance rule that better takes into consideration both upper and lower bounded resources, in particular soft ones, in the solution of the SPPRC with dynamic programming. Moreover, the literature on generic SPPRC, on the NSP and on other related problems provides a large number of algorithmic features that sometimes produced a significant speed-up in the solution of the pricing problem. We have implemented the most promising among them and assessed their impact on a set of INRC instances. Finally, we have also implemented strong-branching, diving heuristic, and a very promising rotation MIP heuristic that produces high quality solutions in parallel of the branch-and-price resolution.

Software design contributions We designed a generic representation of the constraints so that each generic constraint can capture a large set of operational constraints including those described for the INRC competitions. These generic constraints should also be sufficient to take into consideration any type of constraints listed by Cheang et al. [Che+03]. Moreover, considering that other needs may appear in the future, we opted for a code structure that should allow future users to integrate them in the solver. The already developed generic constraints provide many examples to make this task easier.

The algorithms described in this article are implemented in C++ and call only free and open third party libraries if necessary. The resulting code is publicly shared¹ for reproduction of the results, future comparisons, improvements and extensions. To the best of our knowledge, it is the first open-source code that can handle all the hard and soft constraints of the INRC instances.

Computational results Our tests are based on the three benchmarks that were most commonly used in the literature: the Nurse Rostering Problem (NRP)² and INRC-I benchmarks and the static INRC-II instances that have already been considered in the literature. Those benchmarks capture a large diversity of the hard and soft constraints found in the literature. They involve schedules of 8 to 150 nurses over a planning horizon of 2 to 52 weeks, where the nurses can have up to four different skills and each day is divided into up to 32 shifts. Our algorithm is able to solve every instance of INRC-I except one, all static 4-weeks instances of INRC-II, and several 8-weeks instances of INRC-II to optimality. It is also able to find the best known solutions for 19 over 24 NRP instances and it improves two lower bounds with respect to the literature. For some of these results, we had to distribute several

¹The code implementing the methods described in this article is publicly shared on the Git repository <https://github.com/wssuite/NurseScheduler>.

²Instances and best results can be found here <http://www.schedulingbenchmarks.org/nrp/index.html>

tasks on up to eight cores and set a the time limit to one day. Moreover, the numerical tests are an opportunity to report on the experimental comparison of roster and rotation-based formulations on the INRC–II instances.

1.4 Organization of the article

The generic constraints are first detailed in Section 2. We then give the outline of our branch-and-price algorithm in Section 3. In there, we also provide a generic extended formulation of the NSP and describe our branching strategies and the primal heuristics that we developed to improve the upper bounds found by branching. The solution of the pricing problem with a specific dynamic programming algorithm is then described in Section 4. Finally, in Section 5, we report the experimental results of our algorithm, we assess the impact of the acceleration techniques we implemented and carry a comparative study with the best published results.

2 Generic constraints for a general solution of the NSP

In the remainder of the article, we respectively denote as \mathcal{N} , \mathcal{D} , \mathcal{S} and Σ the sets of nurses, days, shifts and skills. The number of days in the planning horizon is denoted as $n_{\mathcal{D}}$, therefore we set $\mathcal{D} = \{1, \dots, n_{\mathcal{D}}\}$. Given that some constraints apply specifically to weekends, we also introduce $n_{\mathcal{W}}$ the number of weekends over the horizon. An *assignment* is then given by any pair (d, s) where $d \in \mathcal{D}$ and $s \in \mathcal{S}$. For a more concise presentation of our models, we include in \mathcal{S} a None shift that corresponds to no assignment on a given day. Stated otherwise, an assignment (d, None) corresponds to a day off whereas (d, s) is a working day, for any $s \neq \text{None}$. We will also refer to None as the *rest shift* and to the others as *work shifts*.

Browsing the literature on the NSP, it quickly appears that each benchmark is based on a specific set of constraints. Cheang et al. [Che+03] have listed the 16 different families of constraints they could find in the articles published before their review. One purpose of this work is to implement an algorithm that will be able to run on most (if not all) instances of the literature provided that a proper parser is implemented. For this, we defined a set of generic constraints that are able to capture most specific constraints that we encountered in the literature.

The first step is in the definition of generic structures for the shifts, days, weekends and patterns, which are most commonly targeted by the constraints.

- A *shift type* is any set of shifts, *i.e.*, any set $S \subseteq \mathcal{S}$. For instance, a shift type may either stand for an actual working shift such as a work period from 9 a.m. to 5 p.m., or it may refer to any working shift, the rest shift or even any afternoon shift.
- A *day type* is any set of days, *i.e.*, any set $D \subseteq \mathcal{D}$. A day type may thus refer to Monday, May 28, 2022, to any Monday, or to any day. In the remainder, we will thus use the days of the week Monday, Tuesday, etc., to refer to the corresponding day types.
- A *weekend type* is given by any sequence of consecutive days of the week, *e.g.*, Saturday-Sunday, Friday-Saturday-Sunday or Saturday-Sunday-Monday.
- Finally, a *pattern* Π is given by tuple of pairs formed of one day type and one shift type, *i.e.*, $\Pi = ((D_1, S_1), \dots, (D_K, S_k))$ where K is the length of Π and $D_k \subseteq \mathcal{D}$ and $S_k \subseteq \mathcal{S}$ for $k = 1, \dots, K$.

Given these structures, we divided the generic constraints into two groups: those that define a feasible individual nurse schedule and its cost, and those that impact the selection of one schedule per nurse. This decomposition is mostly justified by the column generation procedure we have chosen where the master problem includes the second group while the subproblems include the first one. Each generic constraint can then either be hard or soft. In case, it is soft, a penalty is associated to its violation. When a quantity is both upper and lower bounded, the penalties may be different for the upper and lower bounds.

In short, the master problem constraints are given below.

1. *Coverage*: a minimum number of nurses must work on each shift of each day with each skill;
2. *Assignment*: each nurse must be assigned one roster, and they can only be assigned a skill that they possess.

The subproblems generic constraints are more numerous.

3. *Total shift types*: the total number of assignments of each nurse to a given shift type over the planning horizon is bounded. For instance, the total number of working days is usually upper and lower bounded. The total number of night shifts over the horizon may also be bounded with such constraint;

4. *Total worked weekends*: the number of weekends a nurse can work on the planning horizon is bounded;
5. *Consecutive shift types*: the number of consecutive assignments to the same shift type is bounded. This constraint may refer to consecutive working or resting days, to consecutive specific shifts, etc.;
6. *Consecutive worked weekends*: the number of consecutive weekends worked by each nurse is bounded;
7. *Identical weekend*: assignments on a weekend must all be to identical shift types. This can capture that a weekend must be either completely worked or completely rested, and that every assignment on a weekend must be to the exact same shift;
8. *Forbidden pattern*: the schedule of each nurse must not include a given pattern. For instance, it is usually forbidden to work a morning shift after a night shift, and it may be forbidden to rest less than two days after a sequence of night shifts;
9. *Preferences*: each nurse may have several individual wishes for their schedule. These may refer to a day that they want off or to a specific shift that they want to be assigned to on a particular day.

The above constraints allow to model the problems described in the two INRC competitions (which differ in several ways). They also capture the constraints used in the benchmark described by Curtois and Qu [CQ14], which has often been used in the literature. What is more, new constraints can be added to our implementation by defining one corresponding subclass. The condition is that they comply to the framework of our solution algorithm: the master problem needs to be formulated as an ILP and the subproblems as the search for a resource constrained shortest path.

3 Outline of the branch-and-price based on a roster decomposition

3.1 Roster-based column generation

For all $i \in \mathcal{N}$, we denote as Ω_i the set of all feasible rosters for i , where each roster $j \in \Omega_i$ is characterized by a binary vector a^j such that $a_{d,s}^j = 1$ if and only if roster j includes assignment (d, s) for all $(d, s) \in \mathcal{D} \times \mathcal{S}$.

We formulate the NSP as the search for the minimum cost assignment of nurses to rosters under demand constraints. For this, we define three sets of decision variables: for $i \in \mathcal{N}$ and $j \in \Omega_i$, $x_{ij} = 1$ if roster j is assigned to nurse i and 0 otherwise; for $i \in \mathcal{N}$, $p \in \mathcal{D} \times \mathcal{S}$ and $\sigma \in \Sigma$, $y_{i,p,\sigma} = 1$ if nurse i performs skill σ on assignment p ; and for $p \in \mathcal{D} \times \mathcal{S}$ and $\sigma \in \Sigma$, $z_{p\sigma}$ is the demand that is not satisfied on assignment p for skill σ . The vector of roster costs is given by \mathbf{c}_x , while the penalties corresponding to unsatisfied demands and to the assignment to a skill that a nurse does not possess are respectively denoted as \mathbf{c}_z and \mathbf{c}_y . The latter costs are only consistent with soft constraints, but it is more convenient to treat hard and soft constraints in a unified way by assuming a large penalty when a constraint is hard. Also, the skill assignment cost c_y is zero when it corresponds to a skill that a nurse can perform. As a result, we may formulate the NSP as the following extended integer program, where $b_{p,\sigma}$ is the minimum demand on assignment p for skill σ . We also indicate the dual variables associated to assignment and demand constraints between brackets.

$$\min \quad \mathbf{c}_x^T \mathbf{x} + \mathbf{c}_y^T \mathbf{y} + \mathbf{c}_z^T \mathbf{z} \tag{1a}$$

$$\text{s.t.} \quad \sum_{j \in \Omega_i} x_{ij} = 1, \quad \forall i \in \mathcal{N}, \quad [\alpha_i] \tag{1b}$$

$$\sum_{j \in \Omega_i} a_p^j x_{ij} = \sum_{\sigma \in \Sigma} y_{i,p,\sigma}, \quad \forall i \in \mathcal{N}, \forall p \in \mathcal{D} \times \mathcal{S} \quad [\beta_{i,p}] \tag{1c}$$

$$\sum_{i \in \mathcal{N}} y_{i,p,\sigma} + z_{p\sigma} \geq b_{p,\sigma}, \quad \forall p \in \mathcal{D} \times \mathcal{S}, \forall \sigma \in \Sigma \tag{1d}$$

$$x_{ij}, y_{ip\sigma} \in \{0, 1\}, \quad \forall i, j, p, \sigma \tag{1e}$$

$$\mathbf{w}, \mathbf{z} \geq 0 \tag{1f}$$

Given that it is not possible to enumerate the feasible rosters of a nurse in reasonable time for standard instances, we solve the above formulation by column generation. The algorithm starts with an initial set of rosters $\mathcal{R}_i \subseteq \Omega_i$ for each nurse i . It then considers a *restricted master problem* (RMP) that is obtained from Formulation (1)

by replacing Ω_i by \mathcal{R}_i for all nurses $i \in \mathcal{N}$. We may assume that the linear relaxation of RMP, RMP^{LR} , is feasible given that a feasible solution can always be obtained by adding artificial rosters at prohibitive cost. At each iteration of the column generation procedure, RMP^{LR} is solved to optimality. Using an optimal dual solution of this relaxation, a *pricing problem* is then solved to search for rosters of $\Omega_i \setminus \mathcal{R}_i$ with negative reduced costs. If at least one roster of negative reduced cost is generated, it is added to the restricted formulation; if none is found, the optimal solution of RMP^{LR} is then proved to be optimal for the linear relaxation of Formulation (1).

In the remainder of the section, we describe the pricing problem and sketch its solution algorithm before describing how an integer solution can be obtained by branching or by running a primal heuristic algorithm.

3.2 Pricing rosters by dynamic programming

Construction of individual roster graphs The pricing problem consists in searching for individual rosters with negative reduced costs. Given that different nurses are not linked at this stage, pricing can be done independently for each nurse. For this, we consider an acyclic directed *roster graph* $G_i = (V_i, A_i)$ for each nurse i , where

- V_i includes two artificial origin and sink vertices o and t , and one vertex for each assignment $(d, s) \in \mathcal{D} \times \mathcal{S}$. For each $v \in V_i \setminus \{o, t\}$, we denote as $(\text{day}(v), \text{shift}(v))$, the assignment corresponding to v ;
- A_i includes one arc from the origin to each vertex of the first day, one arc from each vertex of the last day to the sink, and for all days but the last one, one arc from each corresponding vertex to each vertex of the following day. More formally,

$$A_i = \{(o, v) : \text{day}(v) = 1\} \cup \{(v, t) : \text{day}(v) = n_{\mathcal{D}}\} \cup \{(u, v) : \text{day}(v) = \text{day}(u) + 1\}.$$

A small example is given in Figure 1.

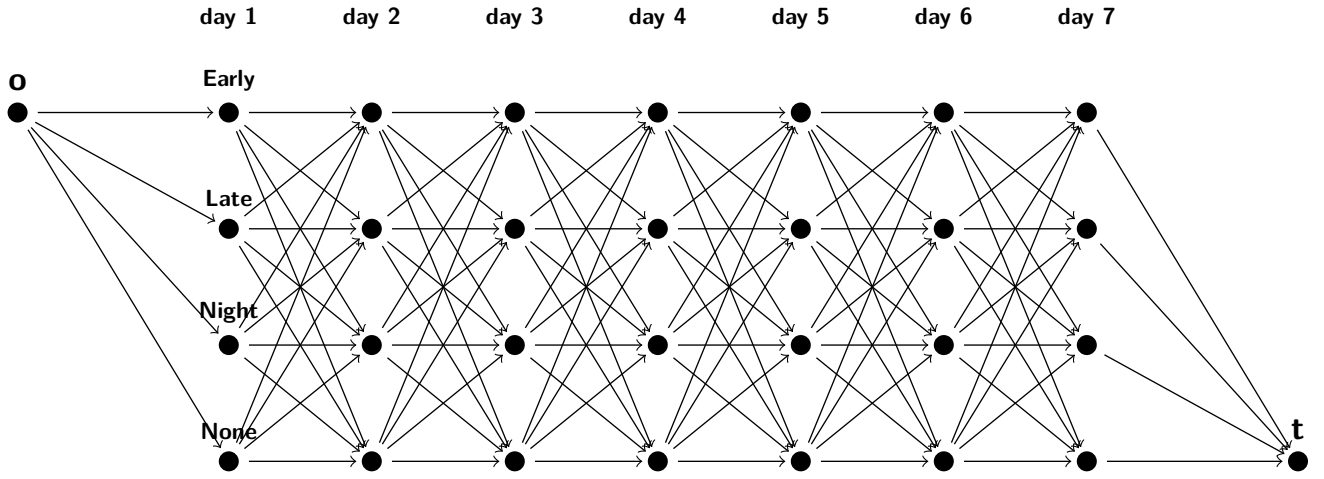


Figure 1: Example of a roster graph network. Here, the horizon includes 7 days and $\mathcal{S} = \{\text{Early}, \text{Late}, \text{Night}, \text{None}\}$.

A roster corresponds to a path from o to t in the roster graph. A feasible roster must also satisfy the individual roster constraints listed in Section 2. We take them into account by either modifying the roster graph or by considering resource constraints. In particular, the constraints on identical weekends, preferences and two-shifts long patterns may all be modeled by modifying the roster graph. For instance, a hard day-off preference is considered by removing all the arcs going to a working shift node on this day, while the same soft preference will lead to adding the corresponding penalty to the cost. Similar modifications can be done on the arcs joining vertices that correspond to different shift types on weekends for the identical weekends constraints.

In contrast, modeling the constraints on total working time/weekends, consecutive shift types/weekends and long forbidden patterns cannot be done without adding a large number of vertices and arcs to the graph. This has been done for instance in [LOR20] in a rotation-based decomposition. But in their approach, they could enumerate the rotations whose lengths were smaller than the minimum number of worked shifts, because this only included rotations with length smaller than three. For rosters, the lower bound may be close to the number of days of the horizon, which would lead to intractable enumeration. As a consequence, these constraints are modeled with a set of resource constraints denoted as \mathbf{R} . With classical resource constraints, an $o - t$ path is feasible if the total

consumption along the path is at most equal to a given upper bound. Here, the constraints may be soft, lower bounds also need to be considered, and the consumption of resources associated to consecutive shifts and forbidden patterns cannot be represented by a single consumption value on each arc. We will give more details on this matter in Section 4. In the remainder of this section, we give an abstract description of the pricing solution algorithm.

With the above modifications of the roster graph and the resource constraints \mathbf{R} , a feasible path and its cost correspond to a feasible roster and its cost. To get the reduced cost of a roster, we also need to model the dual costs of the restricted master problem constraints. Consider nurse $i \in \mathcal{N}$ and a roster $j \in \Omega_i$ with cost c_j that covers assignments (p_1, \dots, p_l) . Referring to Formulation (1a)–(1f), we compute the reduced cost of j as:

$$\bar{c}_j = c_j - \alpha_i - \sum_{k=1}^l \beta_{i,p_k}. \quad (2)$$

To make sure that the cost of any path in the roster graph is equal to the reduced cost of the corresponding roster, we thus subtract α_i from the costs of the arcs going out from the source and we subtract $\beta_{i,p}$ from the costs of the arcs going out from the vertex corresponding to each assignment $p \in \mathcal{D} \times \mathcal{S}$.

Search for a resource-constrained shortest path Following the above construction of the roster graph G , the search for a negative reduced cost roster is equivalent to the search for a negative cost resource-constrained path in G . Searching for the minimum cost feasible path, we thus get a negative reduced-cost roster for nurse i or prove that none exists. Most state-of-the-art approaches for the SPPRC are based on a dynamic programming approach. The algorithm starts with the partial path containing only vertex o before extending it along each arc outgoing from o . The algorithm then proceeds by extending the partial paths obtained at the last iteration along the arcs outgoing from their end vertices until t is reached. Instead of recording each partial path, the algorithm only stores its state as given by its costs, resources consumptions and the last arc it traversed. The efficiency of the algorithm is then determined by its ability to keep only a limited number of *efficient* partial paths at each vertex by removing those that are *dominated*. For a more explicit characterization of extension and dominance, we start with some formal definitions.

Definition 1. Let $v \in V_i$ and let $P = (o, v_1, \dots, v_k)$ be a path from o to v_k with length $|P| = k$. We denote as $\gamma^0(P)$ the cost of P . For any resource $r \in \mathbf{R}$, $\gamma^r(P)$ is the consumption of r along P .

Extension Let $(v_k, v_{k+1}) \in A$, the extension of P along $(v_k, v_{k+1}) \in A_i$ is the path $[P, v_{k+1}] = (o, v_1, \dots, v_k, v_{k+1})$.

Completion A completion of P , $[P, \bar{P}]$, is a path from o to t . The path \bar{P} is then called a suffix of P .

Feasibility Path P is feasible if and only every hard resource constraint is satisfied along P .

Dominance Let Q be a second path from o to v . Path P dominates Q if and only if any feasible completion of Q , $[Q, \bar{Q}]$, yields a feasible completion of P , $[P, \bar{Q}]$, such that $\gamma^0([P, \bar{Q}]) \leq \gamma^0([Q, \bar{Q}])$, and if there is a feasible completion of Q , $[Q, \tilde{Q}]$ such that $\gamma^0([P, \tilde{Q}]) < \gamma^0([Q, \tilde{Q}])$.

Equivalence Path Q is equivalent to P if and only if any feasible completion of Q , $[Q, \bar{Q}]$, yields a feasible completion of P , $[P, \bar{Q}]$, such that $\gamma^0([P, \bar{Q}]) = \gamma^0([Q, \bar{Q}])$. The set of equivalent paths from o to v defines an equivalence class.

To illustrate the above definitions, we may consider the classical SPPRC where each resource constraint corresponds to an upper bound on the resource consumption. Denoting as U_r the upper bound of each resource r , P is infeasible if there is r such that $\gamma^r(P) > U_r$. If P is feasible, path P dominates another feasible path Q if and only if $\gamma^0(P) \leq \gamma^0(Q)$, $\gamma^r(P) \leq \gamma^r(Q)$, $\forall r \in \mathbf{R}$ and $\exists r \in \mathbf{R}$, $\gamma^r(P) < \gamma^r(Q)$. Paths P and Q are equivalent if and only if $\gamma^0(Q) = \gamma^0(P)$ and $\gamma^r(Q) = \gamma^r(P)$, $\forall r \in \mathbf{R}$.

From these definitions, it is straightforward to verify that only non-dominated paths need to be considered for extension, and that at most one representative of each equivalence class needs to be extended at each node. This leads to the generic solution algorithm below (Algorithm 1). The algorithm starts with a partial path including only the origin o , and it extends this path along each arc outgoing from o to build the set of partial paths $\mathcal{P}(v)$ that will be extended from neighbor of o . Since the graph is acyclic and structured into layers corresponding to the days of the planning horizon, the algorithm may then iterate over the days of the horizon to extend the non-dominated partial paths of each corresponding vertex. It may be noted though that the identification of all dominated partial paths requires a necessary and sufficient condition that cannot be verified in a short computational time in general. It may thus be useful to consider only a sufficient condition with the risk that some dominated partial paths will be uselessly extended. Section 4 will be devoted to the specialization of the algorithm for the specific constraints of the NSP.

```

initialization:  $\mathcal{P}(o) := \{(o)\}$ ,  $\mathcal{P}(v) := \emptyset, \forall v \in V_i \setminus \{o\}$ .
1 for  $(o, v) \in A_i$  do
2   | if path  $(o, v)$  is feasible then  $\mathcal{P}(v) \leftarrow \{(o, v)\}$ 
3 for  $d \in \mathcal{D}$  do
4   | for  $u \in V_i(d)$  do
5     | Filter  $\mathcal{P}(u)$  to keep only one path per equivalence class
6     | Filter  $\mathcal{P}(u)$  by removing dominated paths
7     | for  $P \in \mathcal{P}(u)$  do
8       |   | for  $(u, v) \in A_i$  do
9         |   |   | Extend  $P$  along  $(u, v)$  to form  $[P, v]$ 
10        |   |   | if  $[P, v]$  is feasible then  $\mathcal{P}(v) \leftarrow \mathcal{P}(v) \cup \{[P, v]\}$ 
11 if  $\mathcal{P}(t) = \emptyset$  then return  $\emptyset$ 
12 else return  $\operatorname{argmin}_{P \in \mathcal{P}(t)} \{\gamma^0(P)\}$ 

```

Algorithm 1: Dynamic programming search for a constrained shortest path in G_i

3.2.1 Branching

Classical branching decisions on columns do not work within a branch-and-price algorithm: forbidden columns at a given node in the tree can be regenerated by the subproblems, which makes this branching decision useless. Instead, branching decisions are performed on the variables of the original problem before the Dantzig-Wolfe decomposition, e.g., the flow on the edges of the roster graph network (see Figure 1). Therefore, a two-phase branching heuristic has been developed. First, we branch on the flow of the edges corresponding to a rest shift for a given nurse and day: respectively, rest and work edges are removed from the graph of each child node. Second, we branch on a subset of shift types for a given nurse and day: in the same spirit, some edges are removed from each graph.

Furthermore, the branching heuristic must select for each phase on which nurse and day (and subset of shifts for the second phase) the algorithm will branch. As the main goal is to increase the lower-bound when searching to prove optimality, the heuristic aims at maximizing the expected increase of the worst lower-bound of the children. Therefore, the variables closer to 0.5 are selected. When strong-branching is enabled, the linear relaxation of a subset of potential (10, i.e., 5 branching decisions) children is evaluated without column-generation, and the branching decision leading to the best expected increase of the worst lower-bound is selected.

Finally, some knowledge on the nature of the problem is introduced in the variable selection process. First, the assignments on weekends are in general more expensive to cover. Second, some nurses tend to have more fractional and/or expensive rosters. We aim at favoring the variables corresponding to these assignments and nurses in the heuristic. The following score is computed as follows for each variable.

$$\min(f_{ids}, 1 - f_{ids}) + \mu \mathbf{W}(d) + \rho \frac{1 + C_i}{1 + \max_{i'}(C_{i'})} * \frac{N_i}{\max_{i'}(N_{i'})},$$

where f_{ids} is the sum of the flows on the edges corresponding to shift types included in subset s for nurse i on day d , $\mathbf{W}(d)$ indicates whether day d is a weekend, $C_i = \sum_{j \in \Omega_i} c_j x_{ij}$ is the cost of the current fractional roster of nurse i , $N_i = \operatorname{Card}(\{j \in \Omega_i : x_{ij} > 0\})$ is the number of variables of nurse i with a positive value, and (μ, ρ) is a pair of parameters to control the weight of each term. To summarize, the score is composed of three parts: the first one favors fractional flows, the second advantages weekend days and the third one favors nurses with many fractional rosters and a high associated cost. We then branch on the edge with maximum score or test those with maximum scores if strong-branching.

3.2.2 Primal heuristics

In order to get good feasible solutions in the process and increase our chance to prune branching nodes with large lower bounds, we implemented three primal heuristic methods. The first two are simple adaptations of the techniques described in [LOR20] for rotation-based decomposition. In short, we used a diving method where we branch on columns and fix to one the values of those that are integer or close to integrality. Such dives are done frequently at the beginning of the branch-and-price and more scarcely as the algorithm proceeds. We also adapted their large neighborhood search (LNS). When this method is selected and optimality has not been proved after half the time limit, the metaheuristic algorithm is executed to improve the upper bound. Starting with an initial feasible solution, the LNS proceeds by fixing most schedules of the solution and *destroying* the remainder. In a *repair* phase, it then solves a smaller problem where only the destroyed schedules need to be built. Several destruction operators

can be considered and the choice of the operator is made by a random *roulette procedure* that gives priority to the operators that produced improvements in past iterations. In the roster-based decomposition, we mostly destroy complete individual rosters, but we also destroy four consecutive weeks of rosters when considering 8-weeks horizons. The number of nurses whose rosters are destroyed is chosen so that the total number of impacted weeks is equal to 48. Given that our roster-based branch-and-price has been able to solve small instances to optimality within a short computational time, this choice guarantees that the LNS will be able to execute a large number of iterations.

In addition to the above two heuristics, we have implemented a modified version of the classical MIP heuristic for column generation. This heuristic naturally searches for an integer solution by solving the restricted master problem obtained with the columns generated after the solution of the relaxation of the master problem. This is usually done at the root node, but it can also be done at any other branching node. However, compatible rosters are difficult to find, and solving the classical MIP may lead to solutions of poor quality (and, therefore greater than the current upper-bound), and even infeasibility. To overcome this challenge, we come back to the rotation-based formulation that allows more flexibility in the combination of rotations to obtain a good feasible solution. From a current solution of the roster-based formulation, all the roster columns are cut into rotation columns and the associated MIP is then solved. In a parallel environment, this MIP can run permanently: as soon as a better solution is obtained, the new upper-bound is passed to the branch-and-price algorithm and a new MIP based on the current solution is then solved. This rotation MIP heuristic leads to the best solutions, as shown in the experiments.

4 Solution of the shortest path with soft constraints

In this section, we provide a detailed description of the solution of the pricing problem. Given that each nurse can be considered independently in the pricing problem, we focus on a particular nurse i with rotation graph $G_i = (V_i, A_i)$. We denote as \mathbf{H} and \mathbf{S} the set of hard and soft individual constraints of nurse i . For a more concise computation of constraints' violations, we also denote as $[x]^+ = \max\{x, 0\}$ the positive part of any $x \in \mathbb{R}$. Finally, for every constraint $r \in \mathbf{R}$, we denote as U^r the corresponding upper bound and as L^r the lower bound (we can set $L^r = 0$ when there is none). If $r \in \mathbf{S}$, we also denote as c_r^U and c_r^L the linear penalties for violations of the upper and lower limits (respectively).

4.1 Representation of the NSP constraints as resource constraints

As already stated in Section 3.2, the constraints on identical weekends, preferences and two-shifts long patterns may all be modeled by modifying the roster graph. We gave some examples of such modification, and the others are quite straightforward to derive. In the remainder, we show how the remaining constraints can be modeled as resource constraints.

In its most general definition, the SPPRC is a variant of the shortest path problem where a set of resources accumulate along paths and whose value is constrained at vertices and arcs of the path. As a consequence, the description of a resource must specify how it accumulates and how its value is constrained. For each resource $r \in \mathbf{R}$, the accumulation of resource is characterized by a set of resource extension functions (REF) $f_{uv}^r : \mathbb{R} \rightarrow \mathbb{R}, \forall (u, v) \in A_i$. To model soft constraints, we similarly define cost extension functions (CEF) $g_{uv}^r : \mathbb{R} \rightarrow \mathbb{R}, \forall (u, v) \in A_i$. For a given path P with end node u , the extension of resources consumption and cost along some arc $(u, v) \in A_i$ can then be computed as:

- $\gamma^r([P, v]) = f_{uv}^r(\gamma^r(P)), \forall r \in \mathbf{R}$ and
- $\gamma^0([P, v]) = \gamma^0(P) + c_{uv} + \sum_{r \in \mathbf{S}} g_{uv}^r(\gamma^r(P))$.

In the remainder of the section, we provide the specific REFs and CEFs for each constraint of the NSP that needs to be modeled as a resource constraint. For the hard constraints, we also mention the constraints that resource consumptions must satisfy.

Total shift types Let r be a constraint on the total number of assignments to a given shift type S . This constraint can be modeled by a classical resource constraint where the consumption of resource along each arc $(u, v) \in A_i$ is given by a constant value that is equal to 1 if $\text{shift}(v) \in S$ and 0 otherwise. At any arc $(u, v) \in A_i$, the resource/cost extensions functions and the hard constraints on a partial path P can then be defined as follows.

- REF: $f_{uv}^r(\gamma) = \begin{cases} \min\{\gamma + 1, U^r\} & \text{if } \text{shift}(v) \in S \\ \gamma & \text{otherwise} \end{cases}$

- CEF: if $r \in \mathbf{S}$, $g_{uv}^r(\gamma) = \begin{cases} c_r^U & \text{if } \text{shift}(v) \in S \text{ and } \gamma = U^r \\ c_r^L [L^r - \gamma]^+ & \text{if } v = t \\ 0 & \text{otherwise} \end{cases}$
- Hard constraints: if $r \in \mathbf{H}$, $\begin{cases} \text{if } \text{shift}(v) \in S : \gamma^r(P) \leq U^r - 1 \\ \text{if } v = t : \gamma^r(P) \geq L^r \end{cases}$

Observe that penalties due to the violation of the upper limit are taken into consideration when extending a partial path along any arc with the above CEFs. This is done for identifying more dominated paths. In contrast, the penalties due to a violation of the lower bound are added only when reaching the sink, because this violation is non-increasing (with the extension of the path).

Consecutive shift types Let r be a constraint on the consecutive number of assignments to shift type S . These constraints are similar to those found in the variant of the SPPRC studied by Bolívar, Lozano, and Medaglia [BLM14] where the resource consumption is reset when so-called *replenishment arcs* are traversed. Here, any arc (u, v) such that $\text{shift}(u) \in S$ and $\text{shift}(v) \notin S$ is similar to a replenishment arc. When traversing such arc, any started sequence of assignments to the shift type S is ended and the resource value is reset to 0. As a consequence, the lower limit must be satisfied every time a replenishment arc is traversed. For soft constraints, violations of the lower bound must also be paid when traversing such arc. We may thus define the resource/cost extensions and the constraints at any arc $(u, v) \in A_i$ as follows.

- REF: $\forall (u, v) \in A_i, f_{uv}^r(\gamma) = \begin{cases} \min\{\gamma + 1, U^r\} & \text{if } \text{shift}(u) \in S \text{ and } \text{shift}(v) \in S \\ 0 & \text{otherwise} \end{cases}$
- CEF: if $r \in \mathbf{S}$, $g_{uv}^r(\gamma) = \begin{cases} c_r^U & \text{if } \text{shift}(v) \in S \text{ and } \gamma = U^r \\ c_r^L [L^r - \gamma]^+ & \text{if } \text{shift}(u) \in S \text{ and } \text{shift}(v) \notin S \\ 0 & \text{otherwise} \end{cases}$
- Hard constraints: if $r \in \mathbf{H}$, $\begin{cases} \text{if } \text{shift}(v) \in S : \gamma^r(P) \leq U^r - 1 \\ \text{if } \text{shift}(u) \in S \text{ and } \text{shift}(v) \notin S : \gamma^r(P) \geq L^r \end{cases}$

Constraints on weekends The only impacting difference between constraints on shift types and constraints on weekends is that the completion of two different partial paths with the same suffix may lead to (slightly) different consumptions over the suffix. To illustrate this, we must assume the weekends of nurse i is at least three days long, *e.g.*, from Friday to Sunday. We then consider two partial paths P and Q from o to the same vertex u such that $\text{shift}(u) = \text{None}$ and $\text{day}(u)$ is the first Saturday of the planning horizon. Assume also that P goes through a rest shift on the preceding Friday whereas Q goes through a work shift. In such case, if r is a constraint on the total number of working weekends, $\gamma^r(P) = 0$ and $\gamma^r(Q) = 1$. Now, if we extend both P and Q along the same arc going to some node v such that $\text{shift}(v) \neq \text{None}$, the consumption along Q will be unchanged whereas that along P will increase by one unit.

As a consequence, the resource value $\gamma^r(P)$ is not sufficient to represent the state of a partial path P with respect to a constraint on worked weekends. For this, we also need a binary resource value $\delta_r \in \{0, 1\}$ indicating whether the current weekend has already been counted as worked in a partial path P (i.e., $\delta_r(P) = 1$) or not (i.e., $\delta_r(P) = 0$). This resource value has common features with a resource with replenishment since it must be reset at each end of a weekend. To characterize the extension of this resource value, we define one new set of REFs, $h_{uv}^r : \mathbb{R} \rightarrow \mathbb{R}, \forall (u, v) \in A_i$, where

$$h_{uv}^r(\delta) = \begin{cases} 1 & \text{if } \text{shift}(v) \neq \text{None} \text{ and } \text{day}(v) \text{ is on a weekend} \\ \delta & \text{if } \text{shift}(v) = \text{None} \text{ and } \text{day}(v) \text{ is on a weekend} \\ 0 & \text{otherwise} \end{cases}$$

When extended along some arc, the value of $\gamma^r(P)$ can increase only if $\delta_r(P) = 0$. The adaptation of the REFs and CEFs of shift type constraints is then straightforward.

Forbidden patterns They are closely related to the forbidden paths considered by Villeneuve and Desaulniers [VD05], but the structure of the roster graph makes it possible to model it with a single upper-bounded resource.

Let $\Pi = [(D_1, S_1), \dots, (D_K, S_K)]$ be a forbidden pattern with length $K \geq 3$, where D_1, \dots, D_K are day types and S_1, \dots, S_K shift types. Any sub-sequence $[(d_1, s_1), \dots, (d_k, s_k)]$, $k \leq K$, such that $d_1 \in D_1, \dots, d_k \in D_k, s_1 \in S_1, \dots, s_k \in S_k$ is called a *subpattern* of Π with length k . Denoting r the index of the constraint corresponding to Π , we model it with one resource value with upper bound $U^r = K$. For any partial path P from o to $u \in V_i$, $\gamma^r(P)$ indicates the length of the longest subpattern of Π ending at u in P .

If D_1, \dots, D_K refer to specific days or to some specific week days (*e.g.*, Monday, Tuesday and Wednesday), the resource extension along an arc (u, v) is quite straightforward. For $k = \gamma^r(P)$, one must simply compare $\text{day}(v)$ with D_{k+1} and $\text{shift}(v)$ with S_{k+1} . However, if D_1, \dots, D_K include every day of the horizon, some sub-sequences of Π may appear several times in Π . For instance, assume that Π forbids any sequence of shifts (Late, Late, Early) on any sequence of days. If $\gamma^r(P) = 2$, then P ends with a sequence of assignments to shifts (Late, Late). Obviously, if we wish to extend P along some arc (u, v) such that $\text{shift}(v) = \text{Early}$, we get $\gamma^r([P, v]) = 3$. The specificity is that if $\text{shift}(v) = \text{Late}$, the consumption is not reset to 0: the longest subpattern of Π ending at v in $[P, v]$ is (Late, Late) so $\gamma^r([P, v]) = 2$.

To describe the REF, we denote as $\Pi_{\leq l}$ the subpattern of Π with length l , for any $l \leq K$. We then build $\Gamma_\Pi(l)$ as the ordered list of the lengths of the subpatterns of Π that match the end of $\Pi_{\leq l}$. For the simplicity of the algorithm computing the extension, the list $\Gamma_\Pi(l)$ ends with 0 for all l and starts with l for all $l < K$, but it does not include K if $l = K$. In the example above where $\Pi = (\text{Late}, \text{Late}, \text{Early})$, we would get $\Gamma_\Pi(1) = [1, 0]$, $\Gamma_\Pi(2) = [2, 1, 0]$ and $\Gamma_\Pi(3) = [0]$. We finally characterize the resource extension over any arc $(u, v) \in A_i$ as follows.

- REF: $f_{uv}^r(\gamma)$ as described in Algorithm 2 which checks, by decreasing size, if any subpattern could be extended.
- CEF: if $r \in \mathbf{S}$, $g_{uv}^r(\gamma) = \begin{cases} c_r^U & \text{if } f_{uv}^r(\gamma) = K \\ 0 & \text{otherwise} \end{cases}$
- Hard constraints: if $r \in \mathbf{H}$, $f_{uv}^r(\gamma^r(P)) \leq U^r$.

```

1 function  $f_{uv}^r(\gamma)$ :
2   for  $k \in \Gamma_\Pi(\gamma)$  do
3     if  $\text{day}(v) \in D_{k+1}$  and  $\text{shift}(v) \in S_{k+1}$  then
4       return  $k + 1$ ;
5   return 0;

```

Algorithm 2: Resource extension of forbidden patterns

4.2 Dominance rules

The extension functions and the constraints given above guarantee that the dynamic programming (Algorithm 1) will return a minimum cost feasible o - t path. The computational time of the algorithm will depend on its capacity to filter dominated paths though. Our aim is to design rules that are fast to verify and that provide sufficient conditions of dominance. For each considered constraint, one specific test needs to be carried out, and the satisfaction of the sufficient condition depends on the result of every test.

In the remainder of the section, we focus on the comparison of two partial paths P and Q from o to some vertex $v \in V_i$. The size of path P is denoted as $|P| (= |Q|)$. When considering a suffix \bar{Q} of Q , the sequence of vertices of \bar{Q} is denoted as $\bar{Q} = (v_1, \dots, v_{|\bar{Q}|+1})$ and we set $v_0 = v$. Observe that the first filtering task executed in Algorithm 1 is to remove equivalent labels until only one is left per equivalence class. It is straightforward to verify that if $\gamma^0(P) = \gamma^0(Q)$ and $\gamma^r(P) = \gamma^r(Q), \forall r \in \mathbf{R}$, then P and Q are equivalent. As a consequence, we assume that either $\gamma^0(P) \neq \gamma^0(Q)$ or there is $r \in \mathbf{R}$ such that $\gamma^r(P) \neq \gamma^r(Q)$.

4.2.1 Dominance rules for hard constraints

An o - t path is feasible only if it satisfies every hard constraint. As a consequence, P dominates Q only if for all $r \in \mathbf{H}$ and for every completion $[Q, \bar{Q}]$:

$$[Q, \bar{Q}] \text{ satisfies } r \implies [P, \bar{Q}] \text{ satisfies } r. \quad (3)$$

If (3) is true, we say that P dominates Q with respect to (w.r.t.) r . The rule that has often been applied for a resource with both upper and lower limits is that the resource consumption should be the same in both partial paths. We can actually do better in some cases by leveraging that for some paths P , any completion of P is guaranteed to satisfy the upper and/or lower limits of some resources. For this, we denote as \overline{D} the number of days after $\text{day}(v)$ in the planning horizon ($\overline{D} = n_{\mathcal{D}} - \text{day}(v)$) and \overline{W} the number of weekends on the planning horizon after $\text{day}(v)$ (excluding the ongoing weekend if $\text{day}(v)$ is on a weekend). The resulting dominance rule chosen for each type of resource constraint is detailed below.

Proposition 1. *Let $r \in \mathbf{H}$ be any hard resource constraint.*

1. *If r is a constraint on total shift types or on consecutive shift types, P dominates Q w.r.t. r if*

$$[\gamma^r(P) + \overline{D} - U^r]^+ \leq [\gamma^r(Q) + \overline{D} - U^r]^+ \quad \text{and} \quad [L^r - \gamma^r(P)]^+ \leq [L^r - \gamma^r(Q)]^+. \quad (4)$$

2. *If r is a constraint on total or consecutive worked weekends, P dominates Q w.r.t. r if*

$$[\gamma^r(P) + [\delta_r(Q) - \delta_r(P)]^+ + \overline{W} - U^r]^+ \leq [\gamma^r(Q) + \overline{W} - U^r]^+ \quad \text{and} \quad [L^r - \gamma^r(P)]^+ \leq [L^r - \gamma^r(Q)]^+. \quad (5)$$

3. *If r is a constraint on a forbidden pattern Π with length three or more, P dominates Q w.r.t. r if*

$$\gamma^r(P) = \gamma^r(Q) \quad \text{or} \quad \gamma^r(P) \in \Gamma_{\Pi}(\gamma^r(Q)). \quad (6)$$

Proof. We assume that there is no violation of r in P nor Q , and we consider a suffix \overline{Q} such that the completion $[Q, \overline{Q}]$ satisfies r . As a consequence, we are interested only in the violations of constraints that can happen in $[P, \overline{Q}]$ and not in $[Q, \overline{Q}]$. In particular, for constraints on consecutive shift types or worked weekends, we are interested only in the ongoing sequence of shift types or worked weekends. To avoid heavy technicalities, we will refer to the length of this complete sequence as $\gamma^r([P, \overline{Q}])$. This allows for a homogeneous presentation with constraints on total shift and weekend types. We proceed likewise for constraints on forbidden patterns, since we are interested only in patterns that are already started in P or Q .

1. Observe first that by definition of \overline{D} , $\gamma^r(P) \leq \gamma^r([P, \overline{Q}]) \leq \gamma^r(P) + \overline{D}$. By definition of the constraints on shift types, we also have

$$\gamma^r(P) \leq \gamma^r(Q) \Leftrightarrow \gamma^r([P, \overline{Q}]) \leq \gamma^r([Q, \overline{Q}]).$$

Looking at the first inequality of (4), we then see that if $\gamma^r(P) \leq \gamma^r(Q)$ then $[\gamma^r(P) + \overline{D} - U^r]^+ \leq [\gamma^r(Q) + \overline{D} - U^r]^+$. If on the contrary, $\gamma^r(P) > \gamma^r(Q)$,

$$[\gamma^r(P) + \overline{D} - U^r]^+ \leq [\gamma^r(Q) + \overline{D} - U^r]^+ \implies \gamma^r(P) + \overline{D} \leq U^r.$$

As a consequence, if $[\gamma^r(P) + \overline{D} - U^r]^+ \leq [\gamma^r(Q) + \overline{D} - U^r]^+$, then $\gamma^r([P, \overline{Q}]) \leq U^r$ or $\gamma^r(P) \leq \gamma^r(Q)$. Regarding the second inequality, we have similarly that

$$\gamma^r(P) \geq \gamma^r(Q) \implies [L^r - \gamma^r(P)]^+ \leq [L^r - \gamma^r(Q)]^+.$$

And if $\gamma^r(P) < \gamma^r(Q)$,

$$[L^r - \gamma^r(P)]^+ \leq [L^r - \gamma^r(Q)]^+ \implies \gamma^r(P) \geq L^r.$$

As a consequence, if $[L^r - \gamma^r(P)]^+ \leq [L^r - \gamma^r(Q)]^+$, $\gamma^r([P, \overline{Q}]) \geq L^r$ or $\gamma^r(P) \geq \gamma^r(Q)$.

We finally get that if both inequalities are satisfied and $L^r \leq \gamma^r([Q, \overline{Q}]) \leq U^r$ then $L^r \leq \gamma^r([P, \overline{Q}]) \leq U^r$.

2. For constraints on weekends, it is not true that $\gamma^r(P) \leq \gamma^r(Q) \Leftrightarrow \gamma^r([P, \overline{Q}]) \leq \gamma^r([Q, \overline{Q}])$. Indeed, an ongoing weekend may have been worked in Q but not in P . If that is the case, we may have $\gamma^r([P, \overline{Q}]) - \gamma^r(P) = \gamma^r([Q, \overline{Q}]) - \gamma^r(Q) + 1$. Observe then that $[\delta_r(Q) - \delta_r(P)]^+ = 1$ if $\delta_r(P) = 0$ and $\delta_r(Q) = 1$, and 0 otherwise. So, the arguments developed for constraints on shift types remain valid if we consider $\gamma^r(P) + [\delta_r(Q) - \delta_r(P)]^+$ instead of $\gamma^r(P)$ when checking the upper bound.

3. We denote as Π_P and Π_Q the subpatterns of Π at the end of P and Q . Recall that $\gamma^r(P)$ and $\gamma^r(Q)$ are the lengths of Π_P and Π_Q . If $\gamma^r(P) = \gamma^r(Q)$, we then know that if Π appears in $[P, \bar{Q}]$, it must appear in $[Q, \bar{Q}]$. If $\gamma^r(P) \in \Gamma_\Pi(\gamma^r(Q))$, by definition of Γ_Π , it exists a sequence of assignments $\tilde{\Pi}$ such that $\Pi_Q = [\tilde{\Pi}, \Pi_P]$. As a consequence, if Π appears in $[P, \bar{Q}]$, it must also appear in $[Q, \bar{Q}]$. □

In our implementation, we have used the sufficient conditions given in Proposition 1 as dominance tests for hard constraints. In the remainder, for any hard constraint $r \in \mathbf{H}$, the satisfaction of the corresponding test is denoted as $P \succ_r Q$.

4.2.2 Dominance rules for soft constraints

Although soft constraints have usually been identified as difficulties for the solution of the SPPRC, they may actually be leveraged to identify more dominated paths and accelerate the dynamic programming. As an illustrative example, we may consider the simple case with only one soft upper bounded resource and two partial paths P and Q , from o to some vertex v , such that $\gamma^0(P) = 10$ and $\gamma^0(Q) = 20$. Assume that the linear penalty associated with the violation of the resource is equal to 5 and that the consumption along each path is equal to $\gamma(P) = 2$ and $\gamma(Q) = 1$. With the classical dominance rule recalled in Section 3.2, it is impossible to decide whether P dominates Q (because $\gamma(P) > \gamma(Q)$) or Q dominates P (because $\gamma^0(Q) > \gamma^0(P)$). Yet, we see that for any completion $[Q, \bar{Q}]$, $\gamma([P, \bar{Q}]) = \gamma([Q, \bar{Q}]) + 1$ which immediately yields $\gamma([P, \bar{P}]) \leq \gamma([Q, \bar{P}]) + 5$. Given that there is no hard constraint, $[P, \bar{Q}]$ and $[Q, \bar{Q}]$ are both feasible, so the above means that P dominates Q . In the remainder of this section, we formalize this idea and generalize it to all the constraints met in the NSP.

For any $r \in \mathbf{S}$, let $G^r(P, \bar{Q})$ and $G^r(Q, \bar{Q})$ be the penalty paid for violations of constraint r along \bar{Q} in $[P, \bar{Q}]$ and $[Q, \bar{Q}]$, respectively. More formally, with $\bar{Q} = (v_1, \dots, v_{|\bar{Q}|+1})$,

$$G^r(P, \bar{Q}) = \sum_{k=0}^{|\bar{Q}|} g_{v_k v_{k+1}}^r (\gamma^r([P, v_1, \dots, v_k])).$$

We then denote $\Delta^r(P, Q, \bar{Q}) = G^r(P, \bar{Q}) - G^r(Q, \bar{Q})$. Given that the costs of the arcs of \bar{Q} will be paid in both $[P, \bar{Q}]$ and $[Q, \bar{Q}]$, it is straightforward to verify that the following condition is necessary and sufficient for P to dominate Q .

Property 1. *Partial path P dominates Q if and only if for every feasible completion $[Q, \bar{Q}]$, $[P, \bar{Q}]$ is feasible and*

$$\gamma^0(Q) > \gamma^0(P) + \sum_{r \in \mathbf{S}} \Delta^r(P, Q, \bar{Q}). \quad (7)$$

In general, though, the above condition can be difficult to verify for arbitrary partial paths. It may for instance be necessary to extend the two partial paths, which we want to avoid. An alternative is to settle with a sufficient condition that will be based only on the state of the two partial paths. For each soft constraint $r \in \mathbf{S}$, we will thus compute an upper bound $\bar{\Delta}^r(P, Q)$ such that for every completion $[Q, \bar{Q}]$, $\Delta^r(P, Q, \bar{Q}) \leq \bar{\Delta}^r(P, Q)$. Assuming that these quantities are available, we obtain the following sufficient condition of dominance.

Property 2. *For all $r \in \mathbf{S}$, let $\bar{\Delta}^r(P, Q) \geq \Delta^r(P, Q, \bar{Q})$ for all completion $[Q, \bar{Q}]$. Then, P dominates Q if*

$$\begin{cases} P \succ_r Q, \forall r \in \mathbf{H}, \\ \gamma^0(Q) > \gamma^0(P) + \sum_{r \in \mathbf{S}} \bar{\Delta}^r(P, Q) \end{cases} \quad (8)$$

In our implementation, we used the sufficient condition given in (8) as dominance rule. If $P \succ_r Q, \forall r \in \mathbf{H}$ and $\gamma^0(Q) = \gamma^0(P) + \sum_{r \in \mathbf{S}} \bar{\Delta}^r(P, Q)$, either P dominates Q or the two paths are equivalent. In such case, we kept only P for extension. The following proposition provides the actual computation of an upper bound $\bar{\Delta}^r(P, Q)$ for each type of constraint met in the NSP.

Proposition 2. *Let $r \in \mathbf{S}$ be a soft resource constraint, and let \bar{Q} be a suffix of Q such that $[P, \bar{Q}]$ and $[Q, \bar{Q}]$ are feasible.*

1. If r is a constraint on the total number of shift types or on consecutive shift types,

$$\Delta^r(P, Q, \bar{Q}) \leq \begin{cases} c_r^L \left([L^r - \gamma^r(P)]^+ - [L^r - \gamma^r(Q)]^+ \right) & \text{if } \gamma^r(P) \leq \gamma^r(Q) \\ c_r^U \left([\gamma^r(P) + \bar{D} - U^r]^+ - [\gamma^r(Q) + \bar{D} - U^r]^+ \right) & \text{if } \gamma^r(P) > \gamma^r(Q) \end{cases} \quad (9)$$

2. If r is a constraint on total or consecutive worked weekends,

$$\Delta^r(P, Q, \bar{Q}) \leq \begin{cases} c_r^L \left([L^r - \gamma^r(P)]^+ - [L^r - \gamma^r(Q)]^+ \right) & \text{if } \gamma^r(P) \leq \gamma^r(Q) \\ c_r^U \left([\gamma^r(P) + 1 - \delta_r(P) + \bar{W} - U^r]^+ - [\gamma^r(Q) + 1 - \delta_r(Q) + \bar{W} - U^r]^+ \right) & \text{if } \gamma^r(P) > \gamma^r(Q) \end{cases} \quad (10)$$

3. If r is a constraint on a forbidden pattern Π with length three or more,

$$\Delta^r(P, Q, \bar{Q}) \leq c_r^U |\Gamma_\Pi(\gamma^r(Q)) \setminus \Gamma_\Pi(\gamma^r(P))|. \quad (11)$$

Proof. We consider a suffix \bar{Q} of Q . Similarly to the proof of Proposition 1, we are interested in counting the penalties that will be paid while completing P with \bar{Q} but not while completing Q . As a consequence, we can focus on the patterns and on the sequences of shift types/weekends that are already started at the end of P and Q , and we will refer to the length of these completed sequences as $\gamma^r([P, \bar{Q}])$ and $\gamma^r([Q, \bar{Q}])$.

1. Let $0 \leq x \leq \bar{D}$ be the consumption of r in \bar{Q} (this quantity is equal in $[Q, \bar{Q}]$ and in $[P, \bar{Q}]$). Note that x is only the consumption of the initial sequence in \bar{Q} for a constraint on consecutive shift types. We have

$$\gamma^r([P, \bar{Q}]) = \gamma^r(P) + x \text{ and } \gamma^r([Q, \bar{Q}]) = \gamma^r(Q) + x.$$

Referring to the arguments developed in the proof of Proposition 1, we can also deduce that

$$\begin{cases} \Delta^r(P, Q, \bar{Q}) \leq c_r^L \left([L^r - \gamma^r(P) - x]^+ - [L^r - \gamma^r(Q) - x]^+ \right) & \text{if } \gamma^r(P) \leq \gamma^r(Q) \\ \Delta^r(P, Q, \bar{Q}) \leq c_r^U \left([\gamma^r(P) + x - U^r]^+ - [\gamma^r(Q) + x - U^r]^+ \right) & \text{if } \gamma^r(P) > \gamma^r(Q) \end{cases} \quad (12)$$

Functions $x \mapsto [L^r - \gamma^r(P) - x]^+ - [L^r - \gamma^r(Q) - x]^+$ and $x \mapsto [\gamma^r(P) + x - U^r]^+ - [\gamma^r(Q) + x - U^r]^+$ are both stepwise-linear; the former reaches its maximum at 0 while the latter reaches its maximum at \bar{D} . Replacing x with these values yields (9).

2. As in Proposition 1, we must compensate for the difference of structure between constraints on shift types and constraints on weekends by replacing \bar{D} with \bar{W} and by counting possible consumption of resource on the ongoing weekend. This directly yields (10).

3. For each element $k > 0$ of $\Gamma_\Pi(\gamma^r(P))$, there is a subpattern of Π with length k at the end of P . Each subpattern of length $k \in \Gamma_\Pi(\gamma^r(P)) \cap \Gamma_\Pi(\gamma^r(Q))$ that will be completed in $[P, \bar{Q}]$ will also be completed in $[Q, \bar{Q}]$; only the subpatterns with lengths in $\Gamma_\Pi(\gamma^r(Q)) \setminus \Gamma_\Pi(\gamma^r(P))$ may induce a penalty. \square

4.3 Implementation of acceleration techniques

Several acceleration techniques have been used in the past to speed-up the solution of the SPPRC [BLM14]. Good results have also been reported for staff scheduling by Dohn and Mason [DM13] and Gérard, Clautiaux, and Sadykov [GCS16] or for the VRP by Costa, Contardo, and Desaulniers [CCD19]. Here we describe the techniques that looked most relevant for the NSP and that we implemented.

4.3.1 Heuristic pricing methods

Decremental state-space relaxation (DSSR) One may observe that in a context of high demand, the total number of assignments of each nurse will tend to be closer to its upper bound than to its lower bound, and the same goes for the number of assignments in each rotation. As a consequence, we propose to relax the constraints on the lower bounds of these constraints when checking for dominance. If the optimal solution has a negative cost, we add the corresponding column to the master problem. Otherwise, we solve the subproblem once more, but with the lower bound constraints. For a more efficient re-optimization, we store the partial paths that would not have been dominated if lower bounds had been considered.

Restrict the number of propagated labels (RNPL) In [BC14], the number of paths that can be extended is restricted to a given number and the extended paths are chosen heuristically. Given that the graph is acyclic, we instead decided to extend the partial paths with minimum costs at each vertex. If no negative cost roster was found with this heuristic, a second phase is run where every partial path is extended.

4.3.2 Computation of lower bounds

Shortest reverse path from the sink to each vertex (SRP) For any path P from o to v , a lower bound on the cost and on the consumption of each resource in any completion of P can quickly be obtained by searching for reverse shortest paths from the sink t to v . Path P can then be deleted if these lower bounds indicate that every completion will have a non-negative reduced cost or that hard resource constraints will not be satisfied. This technique has often been used for the general SPPRC [BLM14; ZW12], and Burke and Curtois [BC14] also mentioned having implemented it to solve INRC-I instances but they did not assess its impact.

In the NSP, a lower bound on a resource’s consumption is of no real use, because it is usually easy to build a completion that does not consume any unit of resource. As a consequence, we only tested this technique to get lower bounds on the cost of a completion. The reverse shortest paths from t can be computed with only one execution of a $\mathcal{O}(|A_i|)$ shortest path algorithm, because the roster graph is acyclic. However, it has to be executed at every solution of the pricing problem, because dual costs change after every solution of the master problem.

Enumeration of subpaths (ENUM) To strengthen these bounds, we also built a graph where the soft costs of some constraints could be integrated in the arc costs. More specifically, for the constraints on consecutive days off that are found in most benchmarks and for those on consecutive shifts that appear in INRC–II, Legrain, Omer, and Rosat [LOR20] describe how the roster graph can be modified to model the soft constraints with additional arc costs. We adapted their approach by adding arcs that enumerate possible sequences of days off and of assignments to the same shift, and then the shortest reverse path from the sink (**ENUM SRP**).

4.3.3 Bidirectional extension (BIDI)

Bi-directional dynamic programming has originally been used to speed-up Dijkstra’s algorithm for the computation of shortest paths, and it has been reported to bring significant improvements to the solution of VRPs with branch-and-price [RS06]. Instead of extending paths only from the origin to the destination, the algorithm also executes a backward extension from the destination and merges these forward and backward partial paths into complete $o - t$ paths when they join. The vertices where these two paths join are chosen with a view to minimizing the total number of paths generated. This technique can be very efficient for VRPs because the pricing problem consists in the search of elementary paths where the number of paths generated can be exponential in the number of arcs in the paths.

We adapted this technique to NSP, because the number of paths generated may also grow quickly with the number of days in a roster due to the presence of lower bounds. This involves the implementation of two methods for each resource constraint: one for backward extension and the other for merging paths. The vertices corresponding to the middle day of the scheduling horizon were chosen to merge forward and backward paths. This appeared to be the most natural choice for the generation of individual schedules.

5 Computational experiments

Our experiments illustrate the capacity of our approach to handle a large diversity of hard and soft constraints. First, we wish to study the impact of all the algorithmic developments that we made in our search for numerical improvements, and which are described in the previous sections. Second, we assess the performance of our algorithm by solving all the Nurse Rostering Problem (NRP) and INRC-I instances and the INRC–II instances that have previously been used in studies of the static NSP. We also report a comparison with the other methods of the literature that were used to solve the same instances.

In most tests, we aim at proving optimality for as many instances as possible, so we implemented parallel computing where relevant: pricing problems, strong branching and the rotation MIP heuristic. All experiments have been run on processors Intel Gold 6148 Skylake @ 2.4 GHz with 32 GB of memory over 8 threads using our developed open-source C++ code, which depends on the latest Coin-OR libraries (BCP 1.4, CLP 1.17.7 and CBC 2.10.8³) and Gurobi 9.5.0.

³<https://github.com/wssuite/coin>

5.1 Description of the test benchmarks

In our experiments, we have used all the INRC-I instances: 30 sprint, 15 medium and 15 long instances, which consider respectively 10, 30 and 50 nurses. They all consider a four weeks horizon, three to four different shifts, and one or two different skills. Given that the INRC-II benchmark was meant for a stochastic version of the problem, we had to restrict our study to the subset of the instances used by the organizers of the competition when evaluating the competitors' algorithms. Those are the same that have already been used in [LOR20; CGS20]. Their names all start with the same pattern nXXXwY where XXX stands for the number of nurses and Y is the number of weeks of the planning horizon. The pattern is then followed by a long sequence identifying files where the demands are written. For a more compact presentation of the results, we only use shorter names to identify them. The correspondence is given in Table 8 of the Appendix. The instances include 30 to 120 nurses over a four or an eight weeks horizon. They consider four different shifts and four different skills. Among other important differences between the two benchmarks:

- the weekend is always on Saturday and Sunday in INRC-II whereas it can include the Friday in INRC-I,
- INRC-I instances include soft constraints on forbidden patterns with length more than two, whereas they are hard and their length is equal to two in INRC-II,
- possessing a required skill may be a soft constraint in INRC-I while it is always hard in INRC-II,
- there is a soft and a hard lower bound on demand in INRC-II whereas there is only one hard equality constraint in INRC-I.

For the reader who plans on working on the INRC-I benchmark, we should also mention that the correct understanding of the original description of the constraints can be challenging. For this, we had to refer to the validator provided by the organizers of the competition and to the results and solutions already shared by other researchers and competitors.

In comparison to the two INRC benchmarks, the NRP instances can be extremely large, certainly larger than what could be needed for detailed and personalized schedules. The benchmark includes 24 instances whose main characteristics are reported in Table 7. The largest ones consider a one year horizon with up to 32 different shifts per day. Another specificity of the benchmark is that most constraints are hard. In particular, all the constraints on total and consecutive shift types are hard. A detailed comparison of the constraints can be found on the website of the benchmark⁴.

5.2 Assessment of pricing solution

The algorithmic features described in Section 4 are all meant to speed-up the solution of the pricing problems. As a consequence, their impact on computational time can be studied by solving only the linear relaxation of the roster-based formulation (1a)–(1f). Still with a view to spending a reasonable amount of computational resources in this assessment, we focused on six INRC-II instances – three with a 4 weeks horizon and three with a 8 weeks horizon – and three INRC-I instances. We have also run the experiments over only one thread. Given that the number of nurses will not impact our conclusions, we chose the following instances: n30w4-1, n35w4-1, n040w4-1, n030w8-1, n035w8-1, n040w8-1, long3, long_late3 and medium_hidden1.

We first assess the impact of the dominance rules described in Section 4.2 by comparing them to the basic dominance rule where a path P may dominate a path Q only if $\gamma^r(P) = \gamma^r(Q)$ for each upper and lower bounded resource r . Table 1 shows that nearly every strategy implemented is able to decrease significantly the computational time except the enumeration of subpaths which requires a heavy preprocessing. Those results confirm that these strategies work well with a basic dominance rule, as already reported in the literature.

We then assess heuristic pricing algorithms, lower bound computations and bidirectional extension by measuring their impact when considered individually with the improved dominance rule. For heuristic pricing and lower bounds, it was possible to code the accelerating techniques for both INRC-I and INRC-II without extra effort due to the similarity of the two benchmarks. However, each new resource constraint requires the implementation of methods for the backward extension of partial paths and for merging forward and backward partial paths, which requires a decent amount of time to be efficiently implemented. Table 2 shows first that the improved dominance rule eliminates many partial paths. This allows to speed-up the whole solution process by factors up to 28 but, unfortunately it all makes most other strategies inefficient. Only the bidirectional approach is still worth to be implemented especially for bigger instances, and will remain activated for the remaining of the tests.

⁴http://www.schedulingbenchmarks.org/nrp/instances1_24.html

Strategy	Default	ENUM	SRP	ENUM SRP	BIDI	DSSR	RNPL=10	DSSR RNPL=5	DSSR RNPL=10
Instance	Time (s)	Improvement							
n030w4-1	465	4%	58%	2%	55%	37%	36%	33%	33%
n035w4-1	385	18%	54%	0%	50%	51%	32%	36%	29%
n040w4-1	482	30%	58%	-1%	55%	28%	11%	43%	22%
n030w8-1	13772	-1%	65%	1%	61%	34%	14%	20%	28%
n035w8-1	20516	-18%	73%	1%	70%	38%	3%	10%	39%
n040w8-1	16027	13%	67%	-1%	57%	43%	7%	32%	33%
long3	217	0%	24%	-1%	10%	61%	3%	51%	14%
long_late3	3385	-1%	13%	-1%	61%	66%	70%	73%	73%
medium_hidden1	4315	1%	18%	1%	61%	49%	76%	72%	83%
Total		5%	48%	0%	53%	45%	28%	41%	39%

Table 1: Computational time improvements for the basic dominance rule

Strategy	Default	ENUM	SRP	ENUM SRP	BIDI	DSSR	RNPL=10	DSSR RNPL=5	DSSR RNPL=10
Instance	Time (s)	Improvement							
n030w4-1	39	-26%	-8%	-13%	8%	-15%	-5%	-5%	-18%
n035w4-1	40	-23%	-20%	-10%	-15%	-5%	-33%	-23%	-33%
n040w4-1	27	-22%	-7%	-11%	-26%	-19%	-52%	-15%	-59%
n030w8-1	544	-21%	-1%	-4%	29%	-2%	-18%	-13%	-5%
n035w8-1	725	-64%	2%	-8%	30%	-6%	-15%	-26%	-22%
n040w8-1	654	-10%	2%	-2%	19%	-20%	-17%	-18%	-36%
long3	83	-1%	-1%	-6%	45%	4%	-118%	-35%	-112%
long_late3	302	0%	-2%	-4%	33%	-10%	-81%	-71%	-86%
medium_hidden1	303	-1%	-7%	-2%	31%	-1%	-61%	-33%	-61%
Total		-27%	-5%	-8%	7%	-11%	-23%	-16%	-29%

Table 2: Computational time improvements for the improved dominance rule

5.3 Assessment of the branching and primal heuristics

The assessment of the branching and primal heuristics is performed over more difficult instances: n080w4-2, n110w4-1, n030w8-1, n035w8-2, n040w8-1, long3, and medium_hidden1. The selection parameters (μ, ρ) have been tuned in preliminary tests which lead us to use the values (0.2, 0.1) in the remaining experiments.

Table 3 presents the average results of the strong branching heuristic, the two primal heuristics diving and the rotation MIP for a computational time limit of two hours. In the evaluation of the rotation MIP, the MIP solver runs permanently in parallel. Table 3 shows clearly that the primal heuristics decrease the upper bounds while strong branching increases the lower bound. At the same time, one can note that the diving heuristic also decreases the lower bound as more time is spent diving, while strong branching increases the upper bound as more time is spent evaluating the branching decisions. However, the rotation MIP heuristic does not impact the quality of the lower bound as all the computation is done in parallel of the tree exploration. All three strategies are used and combined for the following experiments.

Strategy	diving	rotation MIP	strong branching	diving	rotation MIP	strong branching
Instances	LB			UB		
n080w4-2	0.00%	0.01%	0.05%	1.24%	1.39%	0.31%
n110w4-1	0.21%	-0.01%	0.06%	59.48%	59.48%	0.00%
n030w8-1	-0.15%	0.00%	-0.12%	-0.99%	0.74%	-1.23%
n035w8-2	-0.25%	-0.01%	0.03%	54.53%	56.45%	0.77%
n040w8-1	-0.11%	0.03%	0.17%	-1.54%	2.70%	-2.70%
long3	0.00%	0.00%	0.00%	0.41%	0.41%	0.41%
medium_hidden1	0.00%	0.00%	4.70%	13.87%	10.95%	-21.90%
Average	-0.04%	0.00%	0.70%	18.14%	18.88%	-3.48%

Table 3: Improvement of the lower (LB) and upper (UB) bounds compared to the default branch-and-price algorithm

5.4 Comparison to the literature

The assessment of the default algorithm is done on all the INRC-I instances and on the INRC-II instances used by Ceschia, Guido, and Schaerf [CGS20]. This amounts to a total of 100 instances where the number of nurses ranges from 10 to 120 and the schedules must be computed on 4 to 8 weeks. Finally, the default algorithm is also assessed on the NRP instances that are very different as most constraints are hard, the horizon ranges from 2 weeks to a whole year, and there are up to 32 different shifts per day. The purpose of this comparison is to illustrate the capacity of our code to handle several different benchmarks and assess its capacity to solve the instances to optimality or improve on the literature. In our tests, we set the default time limit to one hour per week in the planning horizon, not exceeding a one-day limit.

5.4.1 INRC-I

A review of the methods applied to the INRC-I and their results are summarized in [Abu+21]. The best known upper bounds are listed on the competition website. Many solution algorithms are based on metaheuristics and they do not compute lower bounds. The only reports on lower bounds that we found in the literature are made by [BC14] and Santos et al. [San+16a], but the former did not give any result on the 20 hidden instances.

Type	early				late				hidden			
Instance	Time (s)	root LB	LB	UB	Time (s)	root LB	LB	UB	Time (s)	root LB	LB	UB
sprint1	107	56	56	56	144	37	37	37	29	31.5	32	32
sprint2	62	58	58	58	15	41.4	42	42	11	32	32	32
sprint3	145	51	51	51	36	47.8	48	48	21	62	62	62
sprint4	30	58.5	59	59	107	72.5	73	73	26	66	66	66
sprint5	134	57	58	58	174	43.7	44	44	90	59	59	59
sprint6	22	54	54	54	61	41.5	42	42	31	129.7	130	130
sprint7	26	56	56	56	41	42	42	42	11	153	153	153
sprint8	49	56	56	56	77	17	17	17	36	204	204	204
sprint9	103	55	55	55	87	17	17	17	267	337.3	338	338
sprint10	49	52	52	52	85	42.9	43	43	29	306	306	306
medium1	94	240	240	240	128	156	<u>157</u>	157	–	95.7	<u>103</u>	120
medium2	26	239.2	240	240	94	18	18	18	–	212.4	<u>216</u>	219
medium3	136	235.5	236	236	155	28.2	29	29	341	33.4	34	34
medium4	363	236.2	237	237	58	34.4	35	35	4817	75.5	<u>78</u>	78
medium5	155	302.1	303	303	450	106.7	107	107	329	117.2	<u>118</u>	118
long1	68	197	197	197	360	235	235	235	830	345	<u>346</u>	346
long2	104	218.5	219	219	336	229	229	229	170	88.5	<u>89</u>	89
long3	80	240	240	240	–	218.5	219	220	3408	37.7	<u>38</u>	38
long4	61	303	303	303	138	220.7	221	221	684	21.8	<u>22</u>	22
long5	183	284	284	284	163	82.5	83	83	1012	41	41	41

Table 4: Results for the INRC-I

We report the results of our algorithm in Table 4. As will be the case in the other tables below, the lower and upper bounds displayed with **bold** characters are optimal, the underlined values indicate an improvement with respect to the literature and a dash in the Time column indicates that time limit was reached before optimality could be proved. Overall, these results show that most instances have been proven optimal in a very short computational time (over 8 threads) and that time limit was reached on only three instances. For the two instances medium hidden 1 and 2, optimality was reached during other tests though. More precisely our algorithm was able to prove optimality with other values of μ and ρ (0 and 0.5). However, our algorithm was never able to converge when executed on instance long late 3. Whatever computational time available, the lower bound remains at the root lower bound. One explication could come from the fact that this particular instance has no personalized aspect: there is no request from any nurse for a day/shift on/off. A non-personalized rostering method should be used for this special instance. Finally, the upper-bounds reported in Table 4 were already known, but optimality proofs had not been reported for all of them. More precisely, we see that our algorithm was able to improve the lower bounds of ten instances. Among those, eight allowed to prove the optimality of the best known upper bound.

Instance	Time (s)	root LB	LB	UB
n030w4-1	67	1659.5	1670	1670^H
n030w4-2	377	1809.2	1815	1815^B
n035w4-1	66	1337.1	1360	1360^B
n035w4-2	196	1075.6	1080	1080^H
n040w4-1	161	1535.2	1565	1565^D
n040w4-2	33	1741.7	1750	1750^D
n050w4-1	483	1295.5	1315	1315^D
n050w4-2	175	1302.1	1315	1315^H
n060w4-1	46	2434.9	2450	2450^D
n060w4-2	31	2664.3	2675	2675^D
n070w4-1	37	2370.3	2380	2380^D
n070w4-2	69	2105	2115	2115^H
n080w4-1	2811	3292	3300	3300^D
n080w4-2	7514	3177.1	3185	3185^H
n100w4-1	187	1167.2	1170	1170^B
n100w4-2	39	1777.7	1780	1780^D
n110w4-1	816	2321.3	2330	2330^D
n110w4-2	133	2455	2455	2455^D
n120w4-1	11351	2011.5	2020	2020^H
n120w4-2	70	2045.7	2050	2050^D

Table 5: Results for the static INRC-II over 4 weeks

5.4.2 Static INRC-II

A review of the methods that were applied to this static INRC-II benchmark and a detailed comparison of their results can be found in [LOR20] and [CGS20]. The results of our roster-based approach on all 4-weeks and 8-weeks instances of the INRC-II benchmark are respectively given in Table 5 and 6. The bounds that appear in these tables all improve on those already reported in the literature. To avoid redundancy, we thus omit to underline these values. For a more detailed analysis of the results, the superscripts B, D, and H refer to the moment when the best upper bound has been found: respectively, during the branch-and-price (either at a node or during strong-branching), when diving or when solving the MIP Heuristic.

First and foremost, Table 5 shows that all 4-weeks instances have been solved to optimality, while the algorithms described in the literature had all failed to do so up to now. Among these 20 instances, 18 were solved in less than one hour and all of them were solved in less than three hours.

Although, the 4-weeks instances have all been solved in a reasonable amount of time, a lot more time was generally necessary to prove optimality for the 8-weeks instances (when it could be done). Table 6 shows that for the 8-weeks horizon, 5 instances over 20 could be solved to optimality in less than eight hours. It also allowed to improve all the best known upper and lower bounds of the literature.

Out of curiosity, we also investigated the impact of the time limit on the results by running these tests with a 1-day time limit. This has allowed to close three more instances and improve many upper and lower bounds. For all instances but one, the remaining gap is of the same order as one unit violation of any soft constraint. This may be the sign that the algorithm is not far from solving most other instances to optimality. Most of these upper bounds have been obtained by the MIP heuristic. For further analysis of the primal heuristics, the large neighborhood search (LNS) developed by Legrain, Omer, and Rosat [LOR20] has also been tested with the roster formulation. These results are presented in the column **LNS UB**. They show that the LNS was almost never able to get better performances than the MIP heuristic approach: the LNS obtained a better upper bound only for instance n040w8-2. One may also observe in Table 6 that n110w8-2 was solved optimally in the second run with a 1-day time limit although less than 3 hours were necessary to converge. The reason is that, even though we set the random seed to the same value for every run, the code remains stochastic due to the parallel execution of the MIP heuristic.

We have then conducted the same experiments with CBC instead of Gurobi for the rotation MIP heuristic and Table 9 presents those results in the Appendix. For the 4-weeks instances, all the instances but one have been solved to optimality, but in general with larger computational times. For the 8-weeks instances, four instances (instead of five) have been solved to optimality under 8 hours and all computational times are significantly higher. CBC can thus be a good option for a free and open-source solver as the reported differences may not be impacting in

Instance	Time limit: 8 hours				Time limit: 1 day			
	Time (s)	root LB	LB	UB	Time (s)	LB	UB	LNS UB
n030w8-1	425	1993.7	2010	2010^H				2025
n030w8-2	471	1709.4	1720	1720^D				1740
n035w8-1	-	2407.8	2425	2505 ^H	-	2430	2460 ^H	2475
n035w8-2	-	2152.2	2205	2270 ^H	-	2215	2265 ^H	2265
n040w8-1	-	2463.9	2490	2515 ^H	-	2495	2515 ^B	2605
n040w8-2	-	2284	2305	2330 ^H	-	2310	2325 ^D	2320
n050w8-1	-	4777.5	4805	4810 ^B	-	4805	4810 ^H	4830
n050w8-2	7731	4743.3	4765	4765^B				4775
n060w8-1	-	2099	2120	2150 ^D	-	2125	2145 ^H	2200
n060w8-2	-	2393.4	2410	2435 ^H	-	2415	2430 ^H	2475
n070w8-1	-	4474.7	4505	4530 ^H	69115	4520	4520^B	4555
n070w8-2	-	4636.3	4660	4690 ^H	20312	4675	4675^H	4700
n080w8-1	-	3941.7	3955	4030 ^H	-	3955	4055 ^H	4110
n080w8-2	-	4286.9	4295	4335 ^H	-	4295	4325 ^H	4335
n100w8-1	-	2013.5	2025	2030 ^H	-	2025	2030 ^H	2065
n100w8-2	-	2128.6	2140	2145 ^H	-	2140	2145 ^H	2190
n110w8-1	2456	3990	3990	1315^B				4000
n110w8-2	-	3440	3440	3460 ^H	11430	3440	3440^B	3440
n120w8-1	-	2440	2440	2450 ^D	-	2440	2450 ^B	2450
n120w8-2	14162	2871.7	2875	2875^B				2875

Table 6: Results for the static INRC-II over 8 weeks

practice.

Finally, we have compared these results to the best ones obtained with the rotation approach [LOR20]. The results are reported in Tables 10 and 11 of the appendix. In the tables, we report the improvements in lower and upper bounds that could be achieved by using a roster instead of a rotation formulation. The upper bounds found when using Gurobi for the MIP heuristic are in column UB and those obtain with CBC are in the column CBC UB. To be fair, we should mention that the rotations results were obtained in one day of computation over only one thread, while the roster formulation was solved on 8 threads with a 4 and 8 hours time limit for 4-weeks and 8-weeks instances, respectively. The results show that the roster approach outperforms the rotation approach on all instances with improvements of the lower bounds by more than 4% on average and of the upper bounds by more than 3%. When using CBC though, it can be noticed that the rotation-based formulation finds a better upper bound for five 8-weeks instances. The observation of the increase in the root and final lower bounds also confirm that the most impacting property of the roster-based formulation is that it has a better linear relaxation.

5.4.3 Nurse Rostering Problem benchmark

The last comparison have been made on the 24 instances of the NRP benchmark. Only the largest instances, which cover a whole year, are still open. For several difficult instances though, there exists no publication describing the method that produced the best reported results. For a meaningful comparison, we will thus restrict our analysis to the best published results. The twelve first instances could be solved to optimality within a reasonable computational time (i.e., less than 5 hours) with a commercial mathematical programming solver, running on 16 threads and using up to 128GB RAM (see [Sme18]). However, the larger ones are much more difficult, which causes most approaches to encounter issues during their solution. The main difficulty comes from the fact there are many hard constraints on total shift types. The easier instances do not include any constraint on the maximum number of assignments to specific shifts whereas the most difficult ones may include 6 to 32 of them. If specifically interested by the NRP benchmark, the reader may refer [SQC20] for more details on published results.

The description of the instance, the results of our roster-based approach and the best published lower and upper bound are reported in Table 7. In the best published columns, the superscript indicates the reference where each bound was published. To obtain our results, most accelerating strategies have been activated (RNPL=10, DSSR, SRP, BIDI). They show that our algorithm performs as well as other approaches on this benchmark: it finds the best known upper bounds of 19 instances and proves the optimality of 14 among them. Moreover, our approach was able to solve one open instance and improve one lower and two upper bounds with respect to the literature. Despite

the large time limit, our approach was not able to solve the root linear relaxation of six instances nor compute any feasible solution for the largest four instances. We observed that the roster-based formulation gets stuck in the solution of the subproblems due to the large number of hard constraints. Although the subproblems are solved only with RNPL and DSSR strategies to overcome part of the difficulty, the impossibility to perform soft domination (as described in Section 4.2.2) on the total shift type constraints yield an explosion of the number of nondominated labels. Regarding instances 15 and 19, this did not preclude the MIP heuristic from finding fairly good upper bounds though. Actually, the upper bound found for instance 19 is even far better than the best published one.

instance	weeks	nurses	shifts	best published		our roster-based approach			
				LB	UB	time (s)	root LB	LB	UB
Instance1	2	8	1	607 ¹	607 ¹	2	558	607	607
Instance2	2	14	2	828 ¹	828 ¹	0	828	828	828
Instance3	2	20	3	1001 ¹	1001 ¹	1	1001	1001	1001
Instance4	4	10	2	1716 ¹	1716 ¹	2	1716	1716	1716
Instance5	4	16	2	1143 ¹	1143 ¹	9	1140.6	1143	1143
Instance6	4	18	3	1950 ¹	1950 ¹	9	1949	1950	1950
Instance7	4	20	3	1056 ¹	1056 ¹	112	1054.1	1056	1056
Instance8	4	30	4	1300 ¹	1300 ¹	-	1296.6	1298	1300
Instance9	4	36	4	439 ¹	439 ¹	-	405.7	406	439
Instance10	4	40	5	4631 ¹	4631 ¹	100	4631	4631	4631
Instance11	4	50	6	3443 ¹	3443 ¹	21	3443	3443	3443
Instance12	4	60	10	4040 ¹	4040 ¹	4131	4040	4040	4040
Instance13	4	120	18	1348 ¹	1356 ²	-	1301	1301	<u>1351</u>
Instance14	6	32	4	1278 ¹	1278 ¹	4830	1278	1278	1278
Instance15	6	45	6	3820 ¹	3853 ¹	-	-	-	3879
Instance16	8	20	3	3225 ¹	3225 ¹	63	3223.5	3225	3225
Instance17	8	32	4	5746 ¹	5746 ¹	669	5746	5746	5746
Instance18	12	22	3	4404 ¹	4459 ¹	-	4417.3	<u>4421</u>	4459
Instance19	12	40	5	3144 ¹	3204 ¹	-	-	-	<u>3174</u>
Instance20	26	50	6	4765 ¹	4913 ¹	36089	4769	4769	4769
Instance21	26	100	8	21122 ¹	21402 ²	-	-	-	-
Instance22	52	50	10	-	32126 ²	-	-	-	-
Instance23	52	100	16	-	19704 ²	-	-	-	-
Instance24	52	150	32	-	58480 ²	-	-	-	-

Table 7: Results for the NRP instances: indices 1 and 2 refer to solutions published respectively in [Sme18] and [SQC20]

6 Conclusion

Extensive experiments (weeks of computation) have shown that the proposed improved dominance approach combined with many acceleration techniques has enabled the resolution of large instances to optimality. In particular, all 4-weeks instances of the INRC-II have now optimal solutions while they were open up to now, and optimality was also proven for eight 8-weeks instances while only small gaps remain for the others. Moreover, every INRC-I instance but one has also been closed, and our approach solves to optimality or finds the best upper bounds for most NRP instances. Finally, the proposed solver, when used with CBC, is available as a free and open source software, and still produces very good solutions.

Several improvements could be made to help with the instances that remain open. In particular, we think that instance long_late3 of the INRC-I can be solved to optimality by aggregating similar nurses. Regarding the largest NRP instances, one may raise the purpose of optimizing detailed and personalized schedules over six months or even one year, but if this is of importance, one step towards getting feasible (and hopefully good) solutions could be implement a receding horizon where smaller schedules are iteratively computed. Finally, some other promising avenues could be explored with stochastic counterparts of the nurse rostering problem. With that view, one may observe that the lower bounds computed with the root linear relaxation of the roster-based approach is within a few percents of the optimal value for all instances solved to optimality. This means for instance that in a two-stage

stochastic counterpart where schedules need to be computed as a potential recourse, the linear relaxation of the roster formulation can serve as an excellent approximation of the second-stage problem.

Acknowledgment

For the purpose of Open Access, a CC-BY public copyright licence has been applied by the authors to the present document and will be applied to all subsequent versions up to the Author Accepted Manuscript arising from this submission.



References

- [Abd+21] Mohammed Abdelghany et al. “A Hybrid Variable Neighbourhood Search and Dynamic Programming Approach for the Nurse Rostering Problem”. In: *Journal of Industrial & Management Optimization* 17.4 (2021), p. 2051. DOI: 10.3934/jimo.2020058.
- [Abu+21] Anmar Abuhamdah et al. “A Novel Population-Based Local Search for Nurse Rostering Problem”. In: *International Journal of Electrical and Computer Engineering (IJECE)* 11.1 (Feb. 1, 2021), p. 471. DOI: 10.11591/ijece.v11i1.pp471-480.
- [AJ14] Khaled S. Abdallah and Jaejin Jang. “An exact solution for vehicle routing problems with semi-hard resource constraints”. In: *Computers & Industrial Engineering* 76 (Oct. 2014), pp. 366–377. ISSN: 03608352. DOI: 10.1016/j.cie.2014.08.011. URL: <https://linkinghub.elsevier.com/retrieve/pii/S0360835214002605>.
- [Bar+98] Cynthia Barnhart et al. “Branch-and-Price: Column Generation for Solving Huge Integer Programs”. In: *Operations Research* 46.3 (1998), pp. 316–329. ISSN: 0030364X. DOI: 10.1287/opre.46.3.316.
- [BC14] Edmund K. Burke and Tim Curtois. “New approaches to nurse rostering benchmark instances”. In: *European Journal of Operational Research* 237.1 (2014), pp. 71–81. ISSN: 0377-2217. DOI: 10.1016/j.ejor.2014.01.039.
- [BC89] J. E. Beasley and N. Christofides. “An Algorithm for the Resource Constrained Shortest Path Problem”. In: *Networks* 19.4 (1989), pp. 379–394. DOI: 10.1002/net.3230190402.
- [BCR14] Andrea Bettinelli, Alberto Ceselli, and Giovanni Righini. “A branch-and-price algorithm for the multi-depot heterogeneous-fleet pickup and delivery problem with soft time windows”. In: *Mathematical Programming Computation* 6.2 (June 2014), pp. 171–197. ISSN: 1867-2949, 1867-2957. DOI: 10.1007/s12532-014-0064-0. (Visited on 02/08/2021).
- [BLM14] Manuel A. Bolívar, Leonardo Lozano, and Andrés L. Medaglia. “Acceleration Strategies for the Weight Constrained Shortest Path Problem with Replenishment”. In: *Optimization Letters* 8.8 (Dec. 2014), pp. 2155–2172. DOI: 10.1007/s11590-014-0742-x.
- [BP05] Jonathan F. Bard and Hadi W. Purnomo. “Preference scheduling for nurses using column generation”. In: *European Journal of Operational Research* 164.2 (2005), pp. 510–534. ISSN: 0377-2217. DOI: 10.1016/j.ejor.2003.06.046.
- [Bur+13] Edmund K. Burke et al. “A Time Predefined Variable Depth Search for Nurse Rostering”. In: *INFORMS Journal on Computing* 25.3 (Aug. 2013), pp. 411–419. ISSN: 1091-9856, 1526-5528. DOI: 10.1287/ijoc.1120.0510. URL: <http://pubsonline.informs.org/doi/abs/10.1287/ijoc.1120.0510> (visited on 05/24/2022).
- [CCD19] Luciano Costa, Claudio Contardo, and Guy Desaulniers. “Exact Branch-Price-and-Cut Algorithms for Vehicle Routing”. en. In: *Transportation Science* 53.4 (July 2019), pp. 946–985. ISSN: 0041-1655, 1526-5447. DOI: 10.1287/trsc.2018.0878. URL: <http://pubsonline.informs.org/doi/10.1287/trsc.2018.0878> (visited on 02/08/2021).
- [Ces+19] Sara Ceschia et al. “The Second International Nurse Rostering Competition”. In: *Annals of Operations Research* 274.1-2 (Mar. 2019), pp. 171–186. DOI: 10.1007/s10479-018-2816-0.

- [CGS20] Sara Ceschia, Rosita Guido, and Andrea Schaerf. “Solving the Static INRC-II Nurse Rostering Problem by Simulated Annealing Based on Large Neighborhoods”. In: *Annals of Operations Research* 288.1 (May 2020), pp. 95–113. DOI: 10.1007/s10479-020-03527-6.
- [Che+03] Brenda Cheang et al. “Nurse rostering problems – a bibliographic survey”. In: *European Journal of Operational Research* 151.3 (2003), pp. 447–460.
- [CQ14] Tim Curtois and Rong Qu. *Computational Results on New Staff Scheduling Benchmark Instances*. Technical report. ASAP Research Group, School of Computer Science, University of Nottingham, 2014, p. 5.
- [DDS06] Guy Desaulniers, Jacques Desrosiers, and Marius M. Solomon. *Column generation*. Vol. 5. Springer Science & Business Media, 2006.
- [DM13] Anders Dohn and Andrew Mason. “Branch-and-Price for Staff Rostering: An Efficient Implementation Using Generic Programming and Nested Column Generation”. In: *European Journal of Operational Research* 230.1 (Oct. 2013), pp. 157–169. DOI: 10.1016/j.ejor.2013.03.018.
- [GCS16] Matthieu Gérard, François Clautiaux, and Ruslan Sadykov. “Column Generation Based Approaches for a Tour Scheduling Problem with a Multi-Skill Heterogeneous Workforce”. In: *European Journal of Operational Research* 252.3 (Aug. 2016), pp. 1019–1030. DOI: 10.1016/j.ejor.2016.01.036.
- [GTS17] Rafael A.M. Gomes, Túlio A.M. Toffolo, and Haroldo Gambini Santos. “Variable neighborhood search accelerated column generation for the nurse rostering problem”. In: *Electronic Notes in Discrete Mathematics* 58 (2017). 4th International Conference on Variable Neighborhood Search, pp. 31–38. ISSN: 1571-0653. DOI: 10.1016/j.endm.2017.03.005.
- [Has+14] Stefaan Haspeslagh et al. “The first international nurse rostering competition 2010”. In: *Annals of Operations Research* 218.1 (2014), pp. 221–236. ISSN: 1572-9338. DOI: 10.1007/s10479-012-1062-0.
- [HQ12] Fang He and Rong Qu. “A constraint programming based column generation approach to nurse rostering problems”. In: *Computers & Operations Research* 39.12 (2012), pp. 3331–3343. DOI: 10.1016/j.cor.2012.04.018.
- [ID+05] Stefan Irnich, Guy Desaulniers, et al. “Shortest path problems with resource constraints”. In: *Column generation*. Springer US, 2005. Chap. 2, pp. 33–65.
- [Jok66] H.C Joksch. “The Shortest Route Problem with Constraints”. In: *Journal of Mathematical Analysis and Applications* 14.2 (1966), pp. 191–197. DOI: 10.1016/0022-247X(66)90020-5.
- [JSV98] Brigitte Jaumard, Frédéric Semet, and Tsevi Vovor. “A generalized linear programming model for nurse scheduling”. In: *European Journal of Operational Research* 107.1 (1998), pp. 1–18. DOI: 10.1016/S0377-2217(97)00330-5.
- [KK04] Niklas Kohl and Stefan E Karisch. “Airline crew rostering: Problem types, modeling, and optimization”. In: *Annals of Operations Research* 127.1-4 (2004), pp. 223–257.
- [Len20] Warner Lensing. “Heuristic Branch-and-Price Algorithms for the Nurse Rostering Problem”. University of Groningen, 2020.
- [LOR20] Antoine Legrain, Jérémy Omer, and Samuel Rosat. “A Rotation-Based Branch-and-Price Approach for the Nurse Scheduling Problem”. In: *Mathematical Programming Computation* 12 (2020), pp. 417–450. DOI: 10.1007/s12532-019-00172-4.
- [LRS11] Federico Liberatore, Giovanni Righini, and Matteo Salani. “A Column Generation Algorithm for the Vehicle Routing Problem with Soft Time Windows”. In: *4OR* 9.1 (Mar. 2011), pp. 49–82. DOI: 10.1007/s10288-010-0136-6.
- [MV10] Broos Maenhout and Mario Vanhoucke. “Branching strategies in a branch-and-price approach for a multiple objective nurse scheduling problem”. In: *Journal of Scheduling* 13.1 (2010), pp. 77–93. ISSN: 1099-1425. DOI: 10.1007/s10951-009-0108-x.
- [QTY10a] Ali Gul Qurashi, Eiichi Taniguchi, and Tadashi Yamada. “Column generation-based heuristics for vehicle routing problem with soft time windows”. In: *Journal of the Eastern Asia Society for Transportation Studies* 8 (2010).
- [QTY10b] Ali Gul Qureshi, Eiichi Taniguchi, and Tadashi Yamada. “Column Generation-based Heuristics for Vehicle Routing Problem with Soft Time Windows”. en. In: *Journal of the Eastern Asia Society of Transportation Studies* 8 (2010), p. 15.

- [Röm16] Michael Römer. “Future demand uncertainty in personnel scheduling: investigating deterministic lookahead policies using optimization and simulation”. In: *Proceedings of 30th European Conference on Modelling and Simulation*. Ed. by Michael Manitz and Oliver Rose. 2016. ISBN: 978-0-9932440-2-5.
- [RS06] Giovanni Righini and Matteo Salani. “Symmetry Helps: Bounded Bi-Directional Dynamic Programming for the Elementary Shortest Path Problem with Resource Constraints”. In: *Discrete Optimization* 3.3 (Sept. 2006), pp. 255–273. DOI: 10.1016/j.disopt.2006.05.007.
- [San+16a] Haroldo G Santos et al. “Integer programming techniques for the nurse rostering problem”. In: *Annals of Operations Research* 239.1 (2016), pp. 225–251.
- [San+16b] Haroldo G. Santos et al. “Integer programming techniques for the nurse rostering problem”. In: *Annals of Operations Research* 239.1 (2016), pp. 225–251. DOI: 10.1007/s10479-014-1594-6.
- [SC86] Thomas R. Sexton and Young-Myung Choi. “Pickup and Delivery of Partial Loads with “Soft” Time Windows”. In: *American Journal of Mathematical and Management Sciences* 6.3-4 (1986), pp. 369–398. DOI: 10.1080/01966324.1986.10737200.
- [Sme18] Pieter Smet. “Constraint reformulation for nurse rostering problems”. In: *Proceedings of the 12th international conference on the practice and theory of automated timetabling*. PATAT. 2018, pp. 69–80.
- [SQC20] Petter Strandmark, Yi Qu, and Timothy Curtois. “First-Order Linear Programming in a Column Generation-Based Heuristic Approach to the Nurse Rostering Problem”. In: *Computers & Operations Research* 120 (Aug. 2020), p. 104945. DOI: 10.1016/j.cor.2020.104945.
- [TGP07] Mariam Tagmouti, Michel Gendreau, and Jean-Yves Potvin. “Arc routing problems with time-dependent service costs”. In: *European Journal of Operational Research* 181 (2007), pp. 30–39. DOI: 10.1016/j.ejor.2006.06.028.
- [VD05] Daniel Villeneuve and Guy Desaulniers. “The Shortest Path Problem with Forbidden Paths”. In: *European Journal of Operational Research* 165.1 (2005), pp. 97–107. DOI: 10.1016/j.ejor.2004.01.032.
- [ZW12] Xiaoyan Zhu and Wilbert E. Wilhelm. “A Three-Stage Approach for the Resource-Constrained Shortest Path as a Sub-Problem in Column Generation”. In: *Computers & Operations Research* 39.2 (Feb. 2012), pp. 164–178. DOI: 10.1016/j.cor.2011.03.008.

A Complete experiment results

$ \mathcal{N} $	$ \mathcal{W} = 4$		$ \mathcal{W} = 8$	
	instance	short name	instance	short name
30	n030w4_1.6-2-9-1	n30w4-1	n030w8_1.2-7-0-9-3-6-0-6	n30w8-1
	n030w4_1.6-7-5-3	n30w4-2	n030w8_1.6-7-5-3-5-6-2-9	n30w8-2
35	n035w4_0.1-7-1-8	n35w4-1	n035w8_0.6-2-9-8-7-7-9-8	n35w8-1
	n035w4_2.8-8-7-5	n35w4-2	n035w8_1.0-8-1-6-1-7-2-0	n35w8-2
40	n040w4_0.2-0-6-1	n40w4-1	n040w8_0.0-6-8-9-2-6-6-4	n40w8-1
	n040w4_2.6-1-0-6	n40w4-2	n040w8_2.5-0-4-8-7-1-7-2	n40w8-2
50	n050w4_0.0-4-8-7	n50w4-1	n050w8_1.1-7-8-5-7-4-1-8	n50w8-1
	n050w4_0.7-2-7-2	n50w4-2	n050w8_1.9-7-5-3-8-8-3-1	n50w8-2
60	n060w4_1.6-1-1-5	n60w4-1	n060w8_0.6-2-9-9-0-8-1-3	n60w8-1
	n060w4_1.9-6-3-8	n60w4-2	n060w8_2.1-0-3-4-0-3-9-1	n60w8-2
70	n070w4_0.3-6-5-1	n70w4-1	n070w8_0.3-3-9-2-3-7-5-2	n70w8-1
	n070w4_0.4-9-6-7	n70w4-2	n070w8_0.9-3-0-7-2-1-1-0	n70w8-2
80	n080w4_2.4-3-3-3	n80w4-1	n080w8_1.4-4-9-9-3-6-0-5	n80w8-1
	n080w4_2.6-0-4-8	n80w4-2	n080w8_2.0-4-0-9-1-9-6-2	n80w8-2
100	n100w4_0.1-1-0-8	n100w4-1	n100w8_0.0-1-7-8-9-1-5-4	n100w8-1
	n100w4_2.0-6-4-6	n100w4-2	n100w8_1.2-4-7-9-3-9-2-8	n100w8-2
110	n110w4_0.1-4-2-8	n110w4-1	n110w8_0.2-1-1-7-2-6-4-7	n110w8-1
	n110w4_0.1-9-3-5	n110w4-2	n110w8_0.3-2-4-9-4-1-3-7	n110w8-2
120	n120w4_1.4-6-2-6	n120w4-1	n120w8_0.0-9-9-4-5-1-0-3	n120w8-1
	n120w4_1.5-6-9-8	n120w4-2	n120w8_1.7-2-6-4-5-2-0-2	n120w8-2

Table 8: Short names of the INRC-II instances

Instance	Time (s)	LB	UB	Instance	Time (s)	LB	UB
n030w4-1	86	1670	1670	n030w8-1	918	2010	2010
n030w4-2	1425	1815	1815	n030w8-2	4774	1720	1720
n035w4-1	71	1360	1360	n035w8-1	28802	2425	2585
n035w4-2	81	1080	1080	n035w8-2	28824	2195	2365
n040w4-1	170	1565	1565	n040w8-1	28800	2490	2565
n040w4-2	36	1750	1750	n040w8-2	28847	2305	2435
n050w4-1	1076	1315	1315	n050w8-1	28804	4800	4835
n050w4-2	206	1315	1315	n050w8-2	10627	4765	4765
n060w4-1	42	2450	2450	n060w8-1	28825	2115	2275
n060w4-2	31	2675	2675	n060w8-2	28817	2410	2595
n070w4-1	37	2380	2380	n070w8-1	28802	4505	4540
n070w4-2	75	2115	2115	n070w8-2	28803	4660	4730
n080w4-1	7806	3300	3300	n080w8-1	28861	3950	4105
n080w4-2	14400	3185	3190	n080w8-2	28844	4295	4440
n100w4-1	138	1170	1170	n100w8-1	28866	2025	2080
n100w4-2	83	1780	1780	n100w8-2	28840	2140	2215
n110w4-1	637	2330	2330	n110w8-1	10108	3990	3990
n110w4-2	59	2455	2455	n110w8-2	28800	3440	3490
n120w4-1	14288	2020	2020	n120w8-1	28820	2440	2450
n120w4-2	67	2050	2050	n120w8-2	28819	2875	2895

Table 9: Results for the INRC-II using COIN-OR CBC solver

Instance	Rotations		Rosters (over 4 hours)							
	LB	UB	Root LB	Increase	LB	Increase	UB	Decrease	CBC UB	Decrease
n030w4-1	1615	1685	1659.5	2.76%	1670	3.41%	1670	0.89%	1670	0.89%
n030w4-2	1740	1840	1809.2	3.98%	1815	4.31%	1815	1.36%	1815	1.36%
n035w4-1	1250	1415	1337.1	6.97%	1360	8.80%	1360	3.89%	1360	3.89%
n035w4-2	1045	1145	1075.6	2.93%	1080	3.35%	1080	5.68%	1080	5.68%
n040w4-1	1335	1640	1535.2	15.00%	1565	17.23%	1565	4.57%	1565	4.57%
n040w4-2	1570	1865	1741.7	10.94%	1750	11.46%	1750	6.17%	1750	6.17%
n050w4-1	1195	1445	1295.5	8.41%	1315	10.04%	1315	9.00%	1315	9.00%
n050w4-2	1200	1405	1302.1	8.51%	1315	9.58%	1315	6.41%	1315	6.41%
n060w4-1	2380	2465	2434.9	2.31%	2450	2.94%	2450	0.61%	2450	0.61%
n060w4-2	2615	2730	2664.3	1.89%	2675	2.29%	2675	2.01%	2675	2.01%
n070w4-1	2280	2430	2370.3	3.96%	2380	4.39%	2380	2.06%	2380	2.06%
n070w4-2	1990	2125	2105	5.78%	2115	6.28%	2115	0.47%	2115	0.47%
n080w4-1	3140	3320	3292	4.84%	3300	5.10%	3300	0.60%	3300	0.60%
n080w4-2	3045	3240	3177.1	4.34%	3185	4.60%	3185	1.70%	3190	1.54%
n100w4-1	1055	1230	1167.2	10.64%	1170	10.90%	1170	4.88%	1170	4.88%
n100w4-2	1470	1855	1777.7	20.93%	1780	21.09%	1780	4.04%	1780	4.04%
n110w4-1	2210	2390	2321.3	5.04%	2330	5.43%	2330	2.51%	2330	2.51%
n110w4-2	2255	2525	2455	8.87%	2455	8.87%	2455	2.77%	2455	2.77%
n120w4-1	1790	2165	2011.5	12.37%	2020	12.85%	2020	6.70%	2020	6.70%
n120w4-2	1820	2220	2045.7	12.40%	2050	12.64%	2050	7.66%	2050	7.66%
Average				7.64%		8.28%		3.70%		3.69%

Table 10: Results for the INRC-II comparing rotation and roster formulations over 4-weeks instances

Instance	Rotations		Rosters (over 8 hours)							
	LB	UB	Root LB	Increase	LB	Increase	UB	Decrease	CBC UB	Decrease
n030w8-1	1920	2070	1993.7	3.84%	2010	4.69%	2010	2.90%	2010	2.90%
n030w8-2	1620	1735	1709.4	5.52%	1720	6.17%	1720	0.86%	1720	0.86%
n035w8-1	2330	2555	2407.8	3.34%	2425	4.08%	2505	1.96%	2585	-1.17%
n035w8-2	2180	2305	2152.2	-1.28%	2205	1.15%	2270	1.52%	2365	-2.60%
n040w8-1	2340	2620	2463.9	5.29%	2490	6.41%	2515	4.01%	2565	2.10%
n040w8-2	2205	2420	2284	3.58%	2305	4.54%	2330	3.72%	2435	-0.62%
n050w8-1	4625	4900	4777.5	3.30%	4805	3.89%	4810	1.84%	4835	1.33%
n050w8-2	4530	4925	4743.3	4.71%	4765	5.19%	4765	3.25%	4765	3.25%
n060w8-1	1970	2345	2099	6.55%	2120	7.61%	2150	8.32%	2275	2.99%
n060w8-2	2260	2590	2393.4	5.90%	2410	6.64%	2435	5.98%	2595	-0.19%
n070w8-1	4400	4595	4474.7	1.70%	4505	2.39%	4530	1.41%	4540	1.20%
n070w8-2	4540	4760	4636.3	2.12%	4660	2.64%	4690	1.47%	4730	0.63%
n080w8-1	3775	4180	3941.7	4.42%	3955	4.77%	4030	3.59%	4105	1.79%
n080w8-2	4125	4450	4286.9	3.92%	4295	4.12%	4335	2.58%	4440	0.22%
n100w8-1	2005	2125	2013.5	0.42%	2025	1.00%	2030	4.47%	2080	2.12%
n100w8-2	2125	2210	2128.6	0.17%	2140	0.71%	2145	2.94%	2215	-0.23%
n110w8-1	3870	4010	3990	3.10%	3990	3.10%	3990	0.50%	3990	0.50%
n110w8-2	3375	3560	3440	1.93%	3440	1.93%	3460	2.81%	3490	1.97%
n120w8-1	2295	2600	2440	6.32%	2440	6.32%	2450	5.77%	2450	5.77%
n120w8-2	2535	3095	2871.7	13.28%	2875	13.41%	2875	7.11%	2895	6.46%
Average				3.91%		4.54%		3.35%		1.46%

Table 11: Results for the INRC-II comparing rotation and roster formulations over 8-weeks instances