# Chapter 8
# Nurse Rostering Problems

## 8.1 Introduction

Personnel scheduling problems arise from various real-world scenarios, including supermarket staff scheduling, call centre staff allocation, police force scheduling, and, the most studied, nurse rostering in hospitals. Due to the demands of quality healthcare, limited resources, and the tight constraints of specific legislation worldwide, the nurse rostering problem (NRP) has received extensive research attention in the last five decades [26].

The NRP consists of assigning a set of nurses with different skills to a set of shifts of different types on each day or timeslot of a scheduling period, satisfying a set of constraints including coverage, legislation, personal preferences, and problem specific requirements. The objective is to minimize the violations of soft constraints (which ideally should be avoided, i.e. nurse preferences) while satisfying hard constraints (which must be satisfied, i.e. all shift demands must be covered during the scheduling period).

As an NP-complete problem [130], nurse rostering presents a research challenge in operational research and meta-heuristics. Algorithms or methods investigated include mathematical programming in operational research, evolutionary algorithms in artificial intelligence and hybrid approaches across different disciplines. Current research in HH on NRP has led to some interesting results based on the extensive research on meta-heuristic algorithms [30].

Along with the extensive NRP research in the last five decades, benchmark NRP datasets have been established, representing a good coverage of real-world problems with different features. These NRP datasets have been applied in HH research; more details can be found in Appendix B.2.

- *UK* dataset: One of the early datasets is derived from a major UK hospital, where a set of 411 pre-processed valid shift patterns / sequences is defined with associated costs calculated based on various constraint violations.
- *INRC2010* datasets: The First International Nurse Rostering Competition [76] (INRC2010) also established a set of benchmark datasets, aiming to bridge the



N. Pillay, R. Qu, *Hyper-Heuristics: Theory and Applications*,
Natural Computing Series, https://doi.org/10.1007/978-3-319-96514-7_8

gaps between theory and practice and promote advances of a range of new approaches.
- *Nottingham* datasets: An NRP web site has been established at the University of Nottingham, providing a collection of a range of NRP problems derived from hospitals worldwide, and the lower-bound solutions reported in the literature.

The majority of current HH approaches for NRP are selection based (Section 8.3), using a diverse set of techniques at the high level to configure a large set of low-level heuristic (*llh*) perturbative operators associated with different acceptance criteria for several well-established NRP benchmarks widely used in the meta-heuristics community. Based on the success of some meta-heuristic algorithms implemented for real-world NRP, more advances in both meta-heuristics and HH might further address the gaps between research and practice for this highly constrained combinatorial optimization problem.

## 8.2 Low-Level Heuristics for Nurse Rostering Problems

Due to the hard constraints in NRP, perturbative *llh* operators are usually defined subject to the fixed coverage requirement of specific shifts on the same day. Swaps of shifts, either consecutive or not, are made between selected nurses on the same day subject to the nurses' skill types. Different methods employed a different subset of the following *llh* and their extensions in the literature [119, 8, 76]:

- *change shift*: change the shift type of a (randomly) selected nurse based on his / her skill type.
- *swap shifts*: between two nurses, swap shifts on *n* consecutive or non-consecutive weekdays or weekends. Nurses may be randomly selected, or heuristically selected based on the number of their conflicts with others, subject to their shift types or skills.
- *move shift*: move the shift of a nurse to another nurse using certain criteria (randomly without considering the costs, or heuristically with the least cost incurred).
- *ruin and recreate*: un-assign and re-assign all shifts of a set of selected nurses in the roster solution randomly or heuristically.

The above list is not exclusive, but presents the most widely used *llh* in HH for NRP. These *llh* have been used with different settings, mostly as simple operators, to gain useful insights into their impact on the performance of different HH approaches. Note that some of the *llh* of smaller size (which make smaller changes to problem solutions) may be redundant given the other larger *llh* (i.e. one larger *llh* operation may be equivalent to the application of several smaller *llh* operations), but both types have been found to contribute to the flexibility of search at different stages of problem solving for instances of different landscapes.

Along with perturbative operators, acceptance criteria have also been investigated and compared in experimental studies [119]. Improving moves are usually accepted, while worsening moves are accepted using various criteria, to strike a balance between exploration and exploitation. Most widely used move acceptance criteria in meta-heuristics are not problem specific, thus can be easily employed in HH.

- *All Move* or *Naive Acceptance*: All neighbourhood solutions by each *llh* are accepted.
- *Only Improving*: Worse neighbourhood solutions are not accepted to encourage exploitation. This can be *first improvement*, which accepts the first better neighbour obtained, or *best improvement* to accept the best among a set of neighbours.
- *Improving or Equal*: Neighbourhood solutions of equal and better quality are accepted.
- *Late Acceptance*: A solution better than the last *n* previously visited solutions is accepted.
- *Simulated Annealing*: Worse solutions are accepted with a probability dependent on the difference from the neighbourhood solutions, and a temperature parameter. The probability gradually decreases, accepting less worse solutions at later stages of the search.
- *Great Deluge*: Worse solutions within a threshold *t* are accepted, and *t* is reduced during the search. Various strategies can be used to gradually reduce the threshold.
- *Adaptive Iteration Limited Threshold Accepting*: This criterion checks a list of recent neighbourhood solutions, and accepts new best solutions found after a number of worse moves, or uses fitness values of recent moves as a threshold for accepting worse solutions.

In HH, a subset of the above perturbative operators, associated with different acceptance criteria, have been used as operator-acceptance *llh* pairs, and selected in selection HH adaptively [17, 119, 8]. In other selection HH approaches, a fixed acceptance criterion is used at the high level, and only perturbative operators are considered and selected as *llh*. In some selection perturbative HH, these simple operators are integrated with acceptance criteria as one combined *llh*, and selected by the high-level heuristics [36, 52, 12, 4].

## 8.3 Selection Hyper-Heuristics for Nurse Rostering Problems

A diverse set of high-level methods have been used in selection HH, mainly on three benchmark datasets detailed in Appendix B.2, namely the UK dataset with 411 shift patterns, the Nottingham benchmark dataset, and the INRC2010 competition dataset. These methods include choice functions, adaptive strategies, Bayesian network and local search algorithms, and provide interesting findings on selecting perturbative *llh*. A summary of different selection HH research in NRP is presented in Table 8.1; details explained in this section.

Early research in HH developed various approaches for 52 instances derived from a major UK hospital by selecting different perturbative *llh* as listed in Section 8.2.

**Table 8.1** Selection perturbative hyper-heuristics for nurse rostering problems

| | **High-level method** | ***llh*** | **Dataset** |
|---|---|---|---|
| Cowling et al. [52] | Choice Function | 9 perturbative operators | UK |
| Burke et al. [36] | Tabu Search | 9 perturbative operators | UK |
| Aickelin et al. [4] | Bayesian Network | rules to select the 411 shift patterns | UK UK |
| Bai et al. [12] | Simulated Annealing | 9 perturbative operators | UK |
| Bilgin et al. [17] | Random, Choice Function, Dynamic Strategy | 12 *swap shifts* and *move shift* operators, with 4 acceptance criteria | INRC2010 |
| Misir et al. [119] | Two adaptive strategies | 29 *swap shifts* and *move shift* operators, with 7 acceptance criteria | INRC2010 instances |
| Shahriar et al. [173] | Iterated Local Search based on tensor analysis | 4 types of perturbative operators | Nottingham |

A set of 411 pre-defined valid shift patterns has been obtained for this problem, and the *llh* selected by the high-level methods operate upon these shift patterns. The problem solutions thus are improved using *llh* to perturb the available patterns / sequences of shifts rather than individual shifts.

In [52], a choice function learns to select from nine perturbative *llh* with the *first improving* acceptance criterion as defined in Section 8.2 to operate upon the 411 shift patterns in the UK benchmark NRP . The evaluation at the high level therefore rewards each *llh* based on its online performance, i.e. the cost of the resulting roster solutions. This same set of *llh* has also been employed in [36], where *llh* are selected by a high-level tabu search within a unified HH framework for the same UK NRP and also benchmark course timetabling problems. In [12], a simulated annealing HH (SAHH) is hybridized with a genetic algorithm to exploit local optima more efficiently. Based on the performance of the acceptance ratio of the nine *llh*, the 411 shift patterns are selected. Results demonstrate the high efficiency of the SAHH approach compared to those in the literature.

An interesting approach is investigated in [4], where a Bayesian network is used as the high-level method to select a string of rules. Each rule chooses from the collection of 411 shift patterns to build roster solutions. Based on a set of training instances, an estimation of distribution algorithm is used to learn the probabilities of the rules that contribute to constructing high-quality roster solutions. The evaluation at the high level is thus a measure of the likelihood the rules lead to good quality rosters. This novel approach can be seen as a selection constructive HH that selects *llh* rules using a statistical model to combine shift patterns to construct high-quality roster solutions.

For the Nottingham benchmark NRP dataset with a large number of diverse problems at hospitals worldwide, a tensor-based machine learning technique is used in [173] to extract patterns from the performance of *llh*. An iterative multi-stage algorithm is then used to automatically select four types of *llh* operators (mutation, crossover, local search, ruin and recreate) with *improving and equal* and *naive ac-*

*ceptance* in the HyFlex framework [29] (see Appendix A.1) based on the knowledge learned by the machine learning model.

The INRC2010 competition attracted a new line of HH research, where a large number of perturbative *llh* are systematically investigated. In [17], three high-level selection methods, namely random selection, choice function, and dynamic heuristic set strategy, are used to select in total 12 variants of *swap shifts* and *move shift* operators described in Section 8.2. These *llh* are categorized in three groups, by days, weekdays and weekends, and *n* (non)-contiguous days, for two randomly chosen nurses. Results have been compared against solutions obtained from integer linear programming within similar computational time.

Another intensive study in [119] employs a *monitor* to manage two high-level selection methods, namely hill climbing and tournament selection, and seven acceptance criteria to select nine heuristic sets. In total 36 variants of the HH approaches with different settings for the high-level search and *llh* are studied to evaluate its diversification and intensification on 10 instances in INRC2010, and also another two healthcare problems. At the high level, the monitoring mechanism manages at the low level four types of acceptance criteria and in total 29 *llh* perturbative operators of different sizes (number of changes to solutions) and speed (execution speed when applied). The effects of heuristic selection and acceptance are analyzed based on the frequencies of *llh* being called and the number of new best solutions found by each *llh*.

The in-depth analysis in [119, 76] revealed some interesting research findings within selection perturbative HH for NRP. A large number of different perturbative operators and acceptance criteria have been automatically selected and combined to develop general algorithms across different problems. HH can be seen as serving as a general framework, where different elements at the two levels can be configured. It was found that it is crucial to combine *llh* of distinct characteristics that can work together. Such a simple framework can also offer a flexible analytical tool to support effective algorithm design in meta-heuristics.

## 8.4 Discussion

Based on the existing rich research in nurse rostering problems (NRP) using a variety of techniques and algorithms, some interesting findings have been obtained on selection perturbative HH approaches. The analysis conducted on simple *llh* of different behaviours, number of changes (sizes) and execution speed has led to a deeper understanding of their performance on different instances [76]. With developments along different lines of research on problem specific *llh*, classifications of NRP, and diverse sets of benchmark problems, HH can be further extended to address a wider range of research issues in the meta-heuristic communities for these highly constrained combinatorial optimization problems.

In selection perturbative HH, investigations have been conducted to examine the effect of *llh* elements including perturbative / move operators and the associated acceptance criteria, on benchmark NRP [76, 119]. Analysis of the synergies between different operators and acceptance criteria at the low level has provided insights into effective search. In this sense, HH can be seen as an analytical tool and framework, to support in-depth analysis of different elements within local search algorithms. Research findings on the combinations among perturbative *llh*, in conjunction with different acceptance criteria, can be extracted to facilitate advanced design of efficient local search algorithms.

As highly constrained combinatorial optimization problems, various NRP have been investigated with a wide range of constraints and problem specific features from different countries. Compared to the diverse research using a variety of high-level algorithms for vehicle routing and examination timetabling problems (Chapters 7 and 10), current research in HH for NRP has mainly focused on selection perturbative approaches. There is a lack of research on constructive HH for NRP. More developments of constructive HH may require further research findings on effective modelling of the complex NRP problems, hybridized with perturbative approaches. For example, for the UK NRP dataset (Appendix B.2.2), a set of valid shift patterns are pre-defined considering various constraints. Selection HH can thus be effectively applied. For the highly constrained NRP, developing generation HH to generate new effective heuristics presents more challenges compared to selection HH; thus simple heuristics and operators are easier to apply to different instances and problems.

Based on the existing rich literature in NRP, HH can also be advanced by further extending *llh* addressing different constraints in different NRP. Research effort has already been made to categorize *llh*. Classifications of NRP have also been proposed [44], using similar notations to those used in the scheduling literature. Given the highly constrained nature of NRP, problem specific simple elementary *llh* could be classified according to which and how many constraints thy address. Both the HH and meta-heuristics communities may benefit from such a systematic study of the synergies between different categories of constructive and perturbative *llh*.

The extensive study of meta-heuristics for NRP has motivated the establishment of several diverse benchmark datasets; the findings obtained have also motivated HH research. The pre-processing of the dataset from a UK major hospital has demonstrated the effectiveness of constraint handling, by encapsulating the problem complexity into a collection of pre-defined valid shift patterns. Similar studies in the NRP literature [83, 188, 23] have also showed promising results using pre-defined valid shift sequences (also called workstretches or stints) of high quality, i.e. with fewer or no violations of constraints. The other two benchmark datasets, the NRP benchmark site maintained at the University of Nottingham (see more details at http://www.schedulingbenchmarks.org/), and the three tracks of datasets at the first nurse rostering competition INRC2010 provide lower bounds as well as a unified format of the problem description. Such efforts are highly valuable and are strongly encouraged to promote future advances in both HH and meta-heuristics communities.