

Vlerick Repository

New computational results for the nurse scheduling problem: A scatter search algorithm

Authors	Maenhout, Broos;Vanhoucke, Mario
Download date	2025-04-29 17:18:45
Link to Item	http://hdl.handle.net/20.500.12127/1958



the Autonomous Management School of
Ghent University and Katholieke Universiteit Leuven

Vlerick Leuven Gent Working Paper Series 2006/06

**NEW COMPUTATIONAL RESULTS FOR THE NURSE SCHEDULING PROBLEM:
A SCATTER SEARCH ALGORITHM**

BROOS MAENHOUT

MARIO VANHOUCKE

Mario.Vanhoucke@vlerick.be

**NEW COMPUTATIONAL RESULTS FOR THE NURSE SCHEDULING PROBLEM:
A SCATTER SEARCH ALGORITHM**

BROOS MAENHOUT

Ghent University

MARIO VANHOUCKE

Vlerick Leuven Gent Management School

Contact:

Mario Vanhoucke

Vlerick Leuven Gent Management School

Tel: +32 09 210 97 81

Fax: +32 09 210 97 00

Email: Mario.Vanhoucke@vlerick.be

ABSTRACT

In this paper, we present a scatter search algorithm for the well-known nurse scheduling problem (NSP). This problem aims at the construction of roster schedules for nurses taking both hard and soft constraints into account. The objective is to minimize the total preference cost of the nurses and the total penalty cost from violations of the soft constraints. The problem is known to be NP-hard.

The contribution of this paper is threefold. First, we are, to the best of our knowledge, the first to present a scatter search algorithm for the NSP. Second, we investigate two different types of solution combination methods in the scatter search framework, based on four different cost elements. Last, we present detailed computational experiments on a benchmark dataset presented recently, and solve these problem instances under different assumptions. We show that our procedure performs consistently well under many different circumstances, and hence, can be considered as robust against case-specific constraints.

Keywords: meta-heuristics; scatter search; nurse scheduling

1 INTRODUCTION

Personnel scheduling problems are encountered in many application areas, such as public services, call centers, hospitals, and industry in general. For most of these organizations, the ability to have suitably qualified staff on duty at the right time is of critical importance when attempting to satisfy their customers' requirements and is frequently a large determinant of service organization efficiency (Thompson, 1995; Felici and Gentile, 2004). This explains the broad attention given in literature to a great variety of personnel rostering applications (see Ernst et al., 2004a; Ernst et al., 2004b). In general, personnel scheduling is the process of constructing occupation timetables for staff to meet a time-dependent demand for different services while encountering specific workplace agreements and attempting to satisfy individual work preferences. The particular characteristics of different industries result in quite diverse rostering models which leads to the application of very different solution techniques to solve these models. Typically, personnel scheduling problems are highly constrained and complex optimization problems (Ernst et al., 2004b; Glover and McMillan, 1986).

In this paper, a procedure is presented to solve the nurse scheduling problem (NSP) which involves the construction of duty rosters for nursing staff over a pre-defined period. Problem descriptions and models vary drastically and depend on the characteristics and policies of the particular hospital. Due to the huge variety of hard and soft constraints, and the several objective function possibilities, the nurse scheduling problem has a multitude of representations, and hence, many exact and heuristic procedures have been proposed to solve the NSP in various guises. Recent literature surveys (Cheang et al., 2003; Burke et al., 2004) give an overview of all these procedures, and mention simulated annealing, tabu search and genetic algorithms as popular meta-heuristics for the NSP.

In constructing a nurse schedule, nurses need to be assigned to days and shifts in order to meet the minimal coverage constraints and other case-specific constraints and to maximize the quality of the constructed timetable. According to Warner (1976), quantifying preferences in the objective function maintains fairness and guarantees the quality of the nurse roster over the scheduling horizon. The coverage constraints embody the required nurses per shift and per day, and are inherent to any shift scheduling problem. The coverage constraints are handled as soft constraints that can be violated at a certain penalty cost expressed in the objective function. Case-specific constraints (determined by personal time requirements, specific workplace conditions, national legislation, etc...) are handled as hard constraints, for which

no violation is possible whatsoever. The objective is thus to minimize the nurses' preferences, expressed as the aversion to work a particular shift on a particular day, and to obtain a feasible schedule as much as possible subject to different case-specific (hard) constraints. The problem is known to be NP-hard (Osogami and Imai, 2000).

In section 2 of this paper, we briefly review the philosophy of the scatter search template provided by Glover (1998). Moreover, we discuss and illustrate the underlying principles and the implementation of the scatter search framework for the nurse scheduling problem. In section 3, we present new computational results tested on the NSPLib dataset proposed by Vanhoucke and Maenhout (2005). In section 4, conclusions are made and directions for future research are indicated.

2 SCATTER SEARCH FOR NSP

Scatter search (Glover, 1998) is a population-based meta-heuristic in which solutions are combined to yield better solutions using convex linear or non-linear combinations. This evolutionary meta-heuristic differs from other evolutionary approaches, such as genetic algorithms, by providing unifying principles for joining solutions based on generalized path constructions in Euclidian space and by utilizing strategic designs where other approaches resort to randomization. The scatter search methodology is very flexible, since each of its elements can be implemented in a variety of ways and degrees of sophistication. Hence, the scatter search template has been successfully applied in several application areas. However, to the best of our knowledge, the scatter search framework has been applied only once to personnel rostering, more precisely on a labour scheduling problem by Casado et al. (2005). In their paper, they describe the development and implementation of a decision support system for the optimization of the passenger flow by trading off service quality and labour costs at an airport. In their search for the minimal number of employees, their path relinking approach concentrates on the shifts which are staffed differently in the parent solutions.

For an overview of the basic and advanced features of the scatter search meta-heuristic, we refer to Glover and Laguna (2000) and Martí et al. (2006). In the following, we describe our implementation of the scatter search approach to the nurse scheduling problem. The pseudo-code for our generic scatter search template to solve the NSP is written below.

Algorithm Scatter Search NSP

```
Diversification Generation Method
While Stop Criterion not met
    Subset Generation Method
    Subset Combination Method
    Improvement Method
    Reference Set Update Method
Endwhile
```

2.1 The Diversification Generation Method

In this initialization step, a large pool of P initial solution vectors is generated. Since useful information about the structure of optimal solutions is typically contained in a suitably diverse collection of elite solutions, the initial solutions are generated in such a manner a critical level of diversity is guaranteed (Glover and Laguna, 2000). In order to generate a diverse set of initial solutions, we create x solutions using a constructive heuristic and $P - x$ solutions in a random way. The constructive heuristic schedules the nurses in a random sequence taking both preference costs and penalty cost of violating the coverage constraints into account. This greedy heuristic is conceived as a minimum cost flow problem which represents all shifts on all days to which a particular nurse can be assigned to. Since not all case-specific constraints (e.g. the maximum number of assignments) can be modelled in the network, it has been implemented by a k -shortest path approach. Based on this initial population, a subset of the population elements are designated to be reference solutions. This reference set contains b_1 high quality solutions (Refset1) and b_2 diverse solutions (Refset2). The construction of Refset1 starts with the selection of the best b_1 solutions in terms of solution quality out of the P initial solutions. In order to select the diverse solutions (Refset2), the minimum distance between all remaining $P - b_1$ solutions and the b_1 solutions is calculated based on the adjacency degree of Aickelin (1999). In pursuit of diversity, the b_2 solutions with maximal distance will be selected for membership of Refset2.

2.2 The Solution Generation Method

After the initialization phase, scatter search operates on this reference set by combining pairs of reference solutions in a controlled, structured way. Two elements of the reference set are chosen in a systematic way to produce points both inside and outside the convex regions spanned by the reference solutions. Glover and Laguna (2000) suggest to create new solutions out of all two-element subsets. Choosing the two reference solutions out

of the same cluster stimulates intensification, while choosing them from different clusters stimulates diversification. Hence, in our scatter search, the solution method consists of the evaluation of all $b_1 \times b_1$, $b_1 \times b_2$ and $b_2 \times b_2$ combinations in a random sequence.

2.3 The Solution Combination Method

A new solution point is the result of a linear combination of two population elements in the reference set. The process of generating linear combinations of a set of reference solutions may be characterized as generating paths between solutions (Glover and Laguna, 1997; Glover, 1999). A path between solutions will generally yield new solutions that share a significant subset of attributes contained in both parent solutions, which can differ according to the path selected. The moves introduce attributes contributed by a guiding solution and/or reduce the distance between the initiating and the guiding solution. The goal is to capture the assignments that frequently or significantly occur in high quality solutions, and then to introduce some of these compatible assignments into other solutions that are generated by a heuristic combination mechanism.

Our specific combination method relies on problem-specific information of both the initiating and the guiding solution schedule to create a new schedule, and takes four criteria into account. Two of these criteria incorporate objective function related data into account, as follows:

Preference costs: Since the overall objective is to minimize the nurses' aversion towards the constructed work timetable, the day/shift preference cost expressed by each nurse is inherent to any problem instance.

Coverage information: In order to minimize the penalty cost of violating the minimal coverage constraints, the algorithm penalizes those shifts where the coverage constraints are violated. In doing so, the solution combination method biases the initiating solution towards a (more) feasible solution.

The other two criteria incorporate information to maintain the "good" characteristics of both the initiating and the guiding schedule, as follows:

Critical shifts of the initiating solution: In directing the initiating solution towards the guiding solution, the algorithm prevents the removal of critical shifts from the initiating solution, which, in case of removal, would lead to an additional violation of the coverage constraints. In doing so, the algorithm aims at the construction of a new solution point that does not encounter any (additional) violations of the coverage constraints.

Bias to the guiding solution: The algorithm guides the initiating solution to the assignments of the guiding solution, in order to decrease the distance between the two schedules by introducing attributes of the guiding solution

These four elements will be carefully taken into account for each move from an initiating solution to a guiding solution. Since the relinking process of the two solutions out of the reference set can be based on more than one neighbourhood (Glover and Laguna, 2000), our algorithm makes use of two types of neighbourhood moves: a nurse neighbourhood move or a day neighbourhood move.

In the *nurse neighbourhood move*, the schedule of a single nurse of the initiating solution is directed towards the schedule of the corresponding nurse in the guiding solution. Therefore, the algorithm relies on a k-shortest path approach to optimize the schedule for a particular nurse taking into account a weighted average of the four aforementioned elements. The scheduling of a single nurse at minimum cost over the complete scheduling horizon can be considered as a minimum cost flow problem and can be solved by any shortest path algorithm. Moreover, since not all case-specific constraints can be incorporated in a shortest path algorithm, a k-shortest path approach is implemented where the outcome of this algorithm (i.e. a nurse schedule) should be checked whether it is feasible or not with respect to all these constraints. If the outcome is not feasible, a 2nd shortest path will be generated and checked for feasibility. This process continues until the shortest feasible pattern (i.e. the k^{th} shortest path) for the nurse is found. The graph used for our algorithm consists of $\#days * \#shifts$ nodes (plus two extra dummy nodes representing the start and end of the network) representing the daily shift assignments for the nurse under study. An arc (a, b) is drawn to connect node a representing a possible shift assignment on day j to node b representing a shift assignment on day $j + 1$. The distances between nodes are made up of a weighted average of the four abovementioned elements. In calculating the new schedule for a particular nurse, we rely on the algorithm of Martins and Pascoal (2003) which identifies the next shortest path deviating from the previously found shortest paths and is based on Dijkstra's labelling algorithm (Dijkstra, 1959).

In the *day neighbourhood move*, the roster of a single day of the initiating solution is directed towards the roster of the corresponding day of the guiding solution, given the assignments of the nurses on all other days in the initiating solution. To that purpose, we transform a single day roster to a linear assignment problem matrix, and solve it by means of the Hungarian method (Kuhn, 1955). In constructing this matrix, we duplicate each shift column such that each shift has a number of columns that is equal to its coverage requirements. Moreover, we add dummy nurses and/or dummy shifts to allow under- or over-coverage of the coverage requirements. The number of extra dummy nurses equals the total

daily nurse requirements, and penalizes under-coverage when a required shift column has to be assigned to a dummy nurse. The number of extra dummy shifts equals the total number of (non-dummy) nurses and allows over-coverage of the coverage requirements when nurses are assigned to dummy shifts. The cost of assigning non-dummy nurses to one of the dummy shifts is equal to the minimum cost of the (feasible) shifts a nurse can be assigned to. The LAP matrix contains costs associated with the four criteria the solution combination mechanism is based on. Furthermore, the LAP matrix excludes certain shift assignments to cope with the case-specific constraints, taking into account the fixed assignments of all other days of the current solution.

This solution combination method can be best illustrated on an example NSP instance with 5 nurses and a scheduling period of 4 days. We assume that each day consists of three shifts (e.g. early (s_1), day (s_2), night (s_3)) and a free shift (s_4). The nurses' preferences as well as the minimal coverage requirements are displayed in the top table of figure 1. Since “ s_4 ” is used to refer to a free shift, its daily coverage requirements equal zero. We assume some additional case specific constraints as follows: the number of assignments varies between a minimal value of 3 and a maximal value of 4. The consecutive working shifts vary between a minimal value of 2 and a maximal value of 4. The assignment of nurses to maximal one shift per day and the succession constraints are inherent to continuous personnel scheduling. The latter constraint implies forbidden successive assignments between s_3 and s_1 , s_3 and s_2 and s_2 and s_1 . The left top table is assumed to be the initiating solution from *Refset*₂ with a total preference cost of 81 and four coverage violations (the specific assignments have been encircled). Since the algorithm penalizes each coverage violation with a penalty cost of 100, the total solution quality of the initiating solution equals 481. The right table is assumed to be the guiding solution from *Refset*₁ with a total preference cost of 70 and no coverage penalties. The adjacency degree measures the distance between schedules as the sum of zeros (identical day/shift assignment) and ones (different day/shift assignment), which leads to a total distance of 12.

In the remainder of this section, we illustrate the nurse neighbourhood move for nurse 4 (left part) and the day neighbourhood move for day 2 (right part) on our two parent solution schedules.

Insert Figure 1 About Here

Nurse neighbourhood move: The left table below the parent solutions displays the calculations of the four elements of this section, i.e. the nurse’s preference costs, the coverage penalties, the critical shifts and the assigned shifts of the guiding solutions. The coverage penalty is set to 100, while the assigned shifts of the guiding solution have a negative cost of 10. The critical shifts are found for those assigned shifts of the initial schedule of nurse 4 when the difference between the number of scheduled nurses (row ‘Scheduled’) and the assignments of nurse 4 is lower than minimum required number of nurses (row ‘Coverage’).

The sum of all costs in the table results in the cost values on the corresponding arcs of the network representing the scheduling of the nurse over the complete scheduling horizon. The graph counts $4 * 4$ nodes and a start and an end dummy node. The shortest path in the network is $s_3 - s_3 - s_4 - s_1$ with a distance of -291. However, the path is infeasible since the constraint of minimal 2 consecutive working days is violated. Based on the same argument, the 2nd shortest path, i.e. $s_1 - s_3 - s_4 - s_1$ with a distance of -289, is infeasible. The next shortest path is $s_1 - s_1 - s_1 - s_1$ with a distance of -213. This path is feasible for all case-specific constraints and leads to the newly constructed schedule at the left bottom part of the figure. The new solution point has a total solution quality of 386. The total preference cost has increased from 81 to 86, whereas the number of coverage violations has decreased from 4 to 3. The distance between the new initiating solution point and the guiding solution point has also decreased from 12 to 11.

Day neighbourhood move: The right table below the parent solutions displays the calculations of the four elements in a similar way as previously, but from the second day’s point-of-view. The corresponding LAP matrix contains the sum of three of these four elements, i.e. the nurse’s preference costs, the total critical shift cost and the cost of the assigned shifts of the guiding solutions. The coverage penalties have been incorporated implicitly in the structure of the LAP matrix since dummy nurses (dn) have been inserted which penalize the under-coverage of shifts. In contrast, the incorporation of dummy shifts (ds) allows the over-coverage of shifts. The case-specific constraints have been embedded in the LAP matrix by excluding some assignments (denoted by crossed cells). The optimal LAP solution has been encircled and leads to the newly constructed schedule at the right bottom part of the figure. The new solution point has a total solution quality of 374. The total preference cost has decreased from 81 to 74, and the number of coverage violations has also decreased from 4 to 3. The distance between the new initiating solution point and the guiding solution point has not changed and remains 12.

2.4 The Improvement Method

The improvement method applies heuristic processes to improve both total preference cost and coverage infeasibilities of the newly generated solution points. To that purpose, we implemented the three complementary local search algorithms of Maenhout and Vanhoucke (2005), each focusing on a different part of the scheduling matrix, i.e. on the schedule of a single nurse, on a single day for all nurses and on the whole schedule for all nurses.

The *pattern-based local search* aims at the optimization of the schedule for a particular nurse, given the schedules for all other nurses. Scheduling a single nurse over the scheduling horizon can be considered as a minimum cost flow problem to compute a shortest path on a suitably defined graph satisfying the hard constraints and the soft constraints as much as possible. Preference costs and penalty costs of violating the coverage constraints are both taken into account.

The *day-based local search* optimizes the schedule for one day given the assignments of the nurses on all other days. To that purpose, the single day in our schedule is converted to a matrix which can be solved as a linear assignment problem. Moreover, the transformation allows over- and under-coverage the coverage requirements and the matrix consists of both preference and penalty cost information in order to bias the solution of the linear assignment problem towards a coverage feasible shift assignment.

The *schedule-based local search* aims to improve the total preference cost of the schedule by swapping (parts of) schedules between nurses. This problem is solved by defining a linear assignment problem that optimally re-distributes the schedules of the current schedule among the nurses. The re-distribution mechanism has only effect on the total preference cost of the schedule, since the coverage remains unchanged. The algorithm tries first to swap complete schedules and then tries to swap parts of the schedule (i.e. parts of two days, whether or not consecutive) between the nurses.

2.5 The Reference Update Method

After the application of the diversification and intensification process, the child solutions are added to the reference set if certain threshold values for the criteria which evaluate the merit of newly created solution points are met. A newly generated solution may become a member of the reference set either if the new solution has a better objective function than the solution with the worst objective value in $Refset_1$ or if the diversity of the new solution to the reference set is larger than the solution with the smallest distance value in $Refset_2$.

In both cases the new solution replaces the worst and the ranking is updated to identify the new worst solution in terms of either quality or diversity. The reference set is dynamically updated. In contrast to a static update where the reference set is updated after combination of all generated sub-sets, a dynamic update evaluates each possible reference set entrance instantly. In this way, new “best” solutions can be combined faster and inferior solutions are eliminated faster. During the search, diversity in the reference set is maintained through the use of these artificial tiers in the reference set but also through a threshold distance depending on the problem size under study. The latter prevents the duplication of solution points in the reference set and/or the entrance of highly resembling solutions.

3 COMPUTATIONAL RESULTS

In this section, we present computational results for our scatter search procedure tested on the NSPLib problem instances of Vanhoucke and Maenhout (2005). This testset contains 4 sub-sets with 25, 50, 75, and 100 nurses and a 7-days scheduling horizon (this so-called *diverse set* contains $4 * 7290$ instances), and 2 sub-sets with 30 or 60 nurses and a 28-days scheduling horizon (this *realistic set* contains $2 * 960$ instances). The nurses’ preference structure and the coverage requirements of each sub-set are characterized by systematically varied levels of various NSP complexity indicators proposed in Vanhoucke and Maenhout (2005). All sets have been extended by 8 mixes of case-specific constraints which appear frequently in literature (Cheang et al., 2003), i.e. the minimum/maximum number of working assignments, the minimum/maximum number of assignments per shift, the minimum/maximum consecutive working shifts, and the minimum/maximum consecutive working shifts. The testsets and the case-specific constraints can be downloaded from <http://www.projectmanagement.ugent.be/nsp.php>. The tests have been carried out on a Toshiba SPA10 with a 2.4 Ghz processor and 256 Mb RAM, under a stop criterion of 1,000 or 5,000 schedules. In the next sub-section, we compare different neighbourhood combination versions into detail. In section 3.2, we present best known solutions for our large dataset.

3.1 Day Neighbourhood or Nurse Neighbourhood

In order to test the performance of our two solution combination methods, the day neighbourhood (DNH) and nurse neighbourhood (NNH), we have randomly selected 576 instances from the 25- and 288 instances from the 30- nurse instances.

Each neighbourhood move (day or nurse) contains a mix of the 4 elements, i.e. preference cost (PC), coverage penalty (CP), critical shift calculation (CS) and the bias to the guiding solution (BG). Both the nurse and the day neighbourhood moves will be compared with simple and straightforward moves, based on

Random cost (RAN): The cost matrix simply contains random numbers instead of the sum of PC, CP, CS and BG.

Percentage of guiding solution (%GS): A randomly selected part of the day (DNH) or nurse (NNH) of the initiating solution will be replaced by the assignments of the guiding solution.

Complete replacement by the guiding solution (CGS): The day (DNH) or nurse (NNH) of the initiating solution will be completely replaced by the assignments of the guiding solution.

Insert Table 1 About Here

The results displayed in table 1 can be summarized as follows. First, the day-based neighbourhood (DNH) outperforms, on the average, the nurse-based neighbourhood (NNH). Indeed, 13 (10) of the 18 tests results in a DNH cost which is lower than its corresponding NNH cost. However, the best results can be obtained with the NNH method. The best result for the N25 (N30) instances amounts to 305.82 (1,486.36) which outperforms the best known results for the DNH method (306.94 and 1,492.97, respectively). Second, the top 4 results for the NNH have been displayed in bold, and show that the three elements, PC, CS and BG are relevant cost factors (3, 3 and 4 times used, respectively). The CP cost factor has been used only once in the top 4 results, and seems to be less important. Last, the simple CGS approach is the best approach for the DNH method, but is not able to outperform the best known results with the NNH method. The %GS has an average (good) performance for the NNH (DNH) method while the RAN method leads to rather poor results for both solution combination methods.

In section 3.2, we report best known solutions for all the data instances, based on the solution combination method NNH-PC/CS/BG for the diverse set with a 7-days scheduling

horizon (N25, N50, N75 and N100) and the solution method NNH-PC/CP/CS/BG for the realistic set with a 28-days scheduling horizon (N30 and N60).

3.2 New Best Known Solutions

In order to benchmark our results and present best known state-of-the-art solutions, we have tested all 31,080 instances on all case-specific constraint files, resulting in 248,640 instances in total. We have truncated each test after a stop criterion of 1,000 and 5,000 schedules. Table 2 and 3 display the results for respectively the 1,000 and 5,000 schedule stop criterion. The tables display the average solution quality of the tests, split up in the average total preference cost (Avg_Pref) and the average penalty cost (Avg_Pen) which is calculated as the average number of violations of the minimal coverage requirements times the penalty cost of 100, the required CPU time (Avg_CPU), the percentage of files for which a feasible solution has been found (%Feas), and the percentage deviation from the LP based lower bound (%Dev_LP). The latter has been found by a simple and straightforward LP model, and has been used for similar tests by [20]. No LP lower bounds could be provided for the N30 en N60 sets, since the number of constraints exceeded the limits for the industrial LINDO optimization library, version 5.3 (Schrage, 1995).

Insert Table 2 & 3 About Here

In order to fine-tune a number of parameters, we have run our procedure on a small subset of all instances under different parameter settings. In order to find an appropriate balance between the diversification and intensification process, we have combined the three proposed local search heuristics into a variable neighbourhood search. For the N25 instances, all nurses are subject to the pattern-based local search and all days are subject to the day-based local search. The schedule-based local search evaluates possible swaps between whole schedules and two-days sub-schedules. In our tests, 40% of all two-days sub-schedules are swapped between the nurses for all testsets. The population size ($b_1 + b_2$) has been set to 15 (10) for the N25, N50 and N30 (N75, N100 and N60) instances, respectively. Each time, 80% of these population elements (b_1) have been put into the $Refset_1$. Moreover, 40% of all nurses are subject to the NNH move for N25 and N50 instances, 60% for the N75 instances, 20% for the N100 instances and all nurses for the N30 and N60 instances. The total time needed to

perform all our tests accounted for approximately 10 days for the 1,000 schedules tests and 35 days for the 5,000 schedules tests.

The table reveals that the gap between the obtained solutions and the LP based lower bounds is in many cases very small. For some of the cases, the gap is somewhat larger which gives an indication of the constrainedness of the problem instances. We suggest that this table can be used for comparison purposes for future researchers. To that purpose, we would like to call the attention to the strict test design used in our computational results section. Therefore, we rely on a limited number of generated schedules, which is a clear and easily applicable stop criterion that is independent of the computer platform and coding skills (see [11]). The results for each individual instance can be downloaded from www.projectmanagement.ugent.be/nsp.php. Our scatter search algorithm is able to outperform (equalize) 40.5% and 62.06% (27.12% and 1.45%) of respectively the diverse and realistic instances solved by the electromagnetic procedure of [20] with a stop-criterion of 1,000 schedules, and 33.29% and 64.42% (45.02% and 1.74%) of respectively the diverse and realistic instances for the 5,000 schedules. We would like to call researchers to outperform these results and report their solutions on our website.

4 CONCLUSIONS AND FUTURE RESEARCH

In this paper, we have presented a new scatter search procedure for the well-known nurse scheduling problem. To the best of our knowledge, the literature on scatter search for the nurse scheduling problem is completely void. This framework has only been applied once on a similar problem type (labour scheduling) by Casado et al. (2005).

We have investigated the use of two types of solution combination methods, based on the combinations of sub-schedules of nurses or days. Each method calculates the attractiveness of the move based on four criteria. We have shown that the scatter search algorithm leads to promising results and hence might have a bright future in the further development of meta-heuristic optimization algorithms. We have tested our instances based on a generated problem set NSPLib, under a strict test design with a strict stop criterion to facilitate comparison between procedures. All results can be downloaded from our website which can be used to motivate researchers to report new, outperforming state-of-the-art results.

Our main future research intention is as follows. We will aim at the development of hybrid versions of different meta-heuristics, based on knowledge and concepts presented in this and many other research papers. A skilled combination of concepts of different meta-

heuristics can provide a more efficient behaviour and a higher flexibility when dealing with real-world and large-scale problems.

REFERENCES

- Aickelin, U., 1999, "Genetic Algorithms for Multiple-Choice Optimisation Principles", PhD, University of Wales Swansea.
- Burke, E.K., De Causmaecker, P., Vanden Berghe, G. and Van Landeghem, H., 2004, "The state of the art of nurse rostering", *Journal of Scheduling*, 7, 441-499.
- Casado, S., Laguna, M., and Pacheco, J., 2005, "Heuristical Labor Scheduling to Optimize Airport Passenger Flows", *Journal of the Operational Research Society*, 56, 649-658.
- Cheang, B., Li, H., Lim, A., and Rodrigues, B., 2003, "Nurse rostering problems – a bibliographic survey", *European Journal of Operational Research*, 151, 447-460
- Dijkstra, E.W., 1959, "A note on two problems in connexion with graphs", *Numerische Mathematik.*, 1, 269-271.
- Ernst, A.T., Jiang, H., Krishamoorthy, M., Owens, B., and Sier, D., 2004a, "Staff scheduling and rostering: A review of applications, methods and models", *European Journal of Operational Research*, 153, 3-27.
- Ernst, A.T., Jiang, H., Krishamoorthy, M., Owens, B., and Sier, D., 2004b, "An Annotated Bibliography of Personnel Scheduling and Rostering", *Annals of Operations Research*, 127, 21-144.
- Felici, G., and Gentile, C., 2004, "A polyhedral Approach for the Staff Rostering Problem", *Management Science*, 50, 381-393.
- Glover, F., and McMillan, C., 1986, "The General Employee Scheduling Problem: An integration of MS and AI", *Computers and Operations Research*, 13, 563-573.
- Glover, F. and M. Laguna (1997). *Tabu Search*, Kluwer Academic Publishers.
- Glover, F. (1998). "A Template for Scatter Search and Path Relinking," in *Lecture Notes in Computer Science*, 1363, J.K. Hao, E. Lutton, E. Ronald, M. Schoenauer, D. Snyers (Eds.), pp. 13-54.
- Glover F. (1999). "Scatter Search and Path Relinking," In: D Corne, M Dorigo and F Glover (eds.) *New Ideas in Optimisation*, Wiley.

Glover, F., and Laguna, M., 2000, "Fundamentals of Scatter Search and Path Relinking", *Control and Cybernetics*, 3, 653-684

Kuhn, H., 1955, "The Hungarian method for the assignment problem", *Naval Research Logistics*, 2, 83-97.

Maenhout, B. and Vanhoucke, M., 2005, "An Electromagnetism meta-heuristic for the nurse scheduling problem", *working paper 05/316*, Ghent University.

Marti, R., Laguna, M., Glover, F., 2006, "Principles of scatter search", *European Journal of Operational Research*, 169, 359-372.

Martins, E.Q.V., and Pascoal, M.M.B., 2003, "A new implementation of Yen's ranking loopless paths algorithm", *4OR – Quarterly Journal of the Belgian, French and Italian Operations Research Societies*, 1, 121-134.

Osogami, T., and Imai, H., 2000, "Classification of Various Neighbourhood Operations for the Nurse Scheduling Problem", *Lecture Notes in Computer Science*, 1969, 72-83.

Schrage, L., 1995, "LINDO: Optimization software for linear programming", LINDO Systems Inc.: Chicago, IL.

Thompson, G.M., 1995, "Improved implicit optimal modelling of the labour shift scheduling problem", *Management Science*, 43, 595-607.

Vanhoucke, M., and Maenhout, B., "Characterisation and Generation of Nurse Scheduling Problem Instances", *working paper 2005*, Ghent University.

Warner, H.W., 1976, "Scheduling Nursing Personnel According to Nursing Preference: A Mathematical Approach", *Operations Research*, 24, 842-856.

FIGURE 1

The Solution Combination Method in nurse and day neighbourhood space

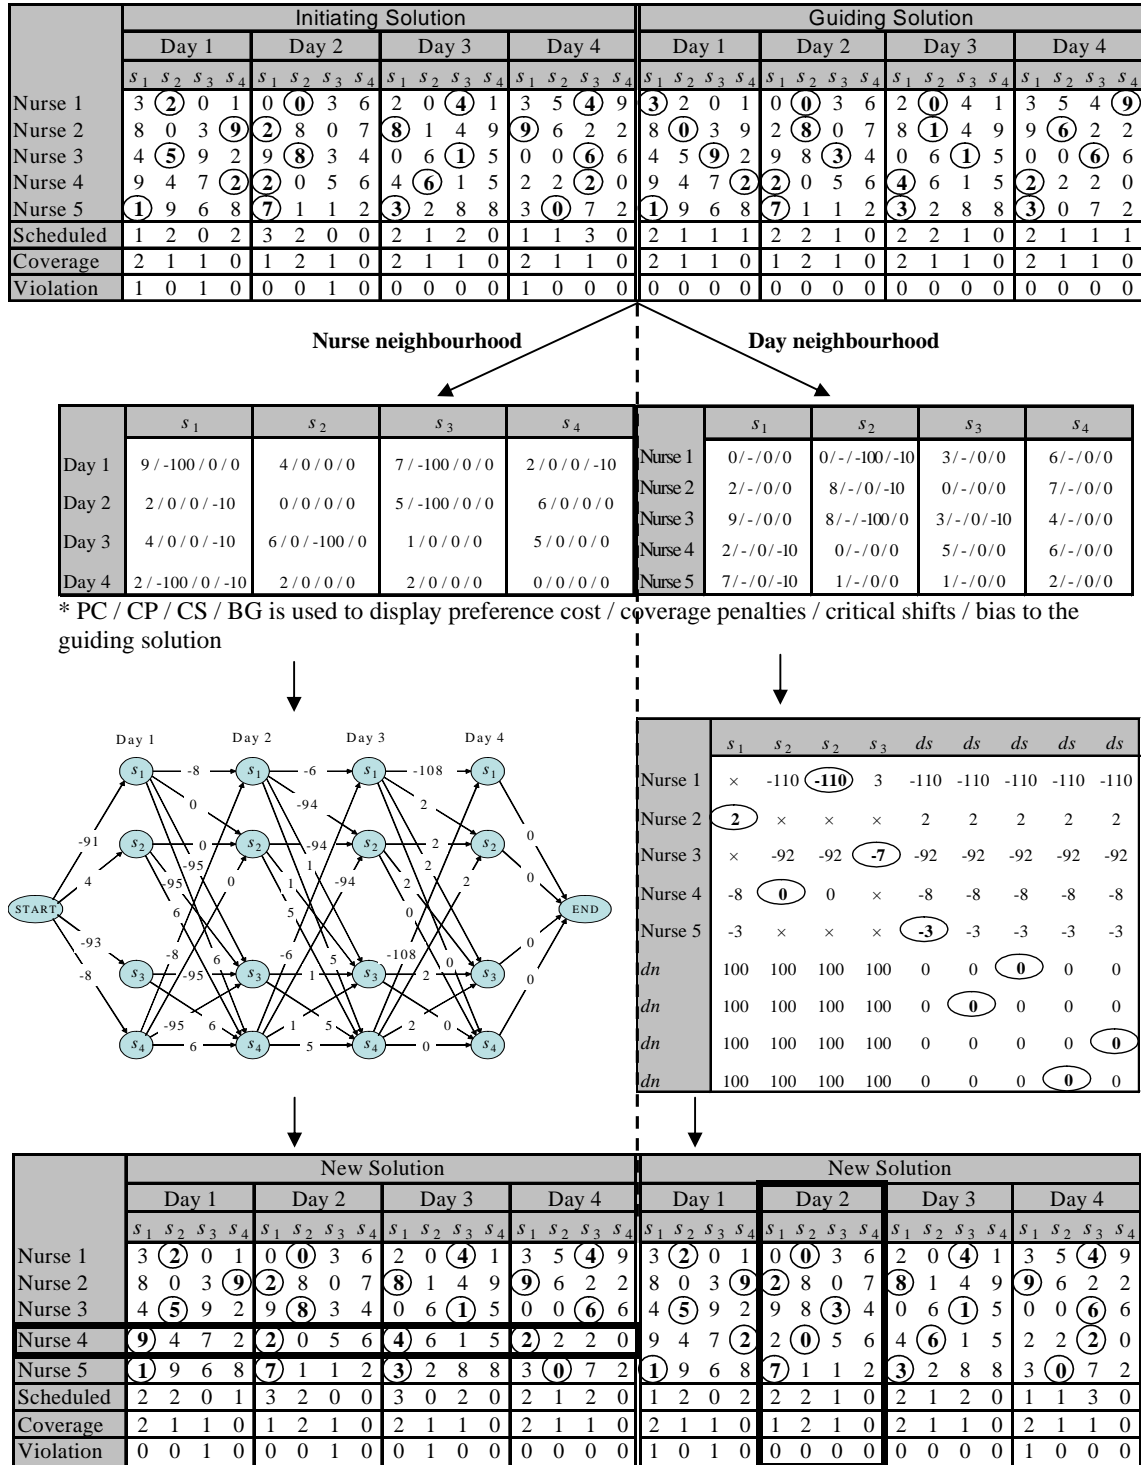


TABLE 1

Average solution quality (Avg_Sol) and Ranking for the various combination methods

	Combination elements				Nurse Neighbourhood				Day Neighbourhood			
	PC	CP	CS	BG	N25		N30		N25		N30	
					Avg_Sol	Ranking	Avg_Sol	Ranking	Avg_Sol	Ranking	Avg_Sol	Ranking
Solution Combination Method	×	-	-	-	310.67	17	1,498.21	9	307.52	14	1,500.20	15
	-	×	-	-	312.26	18	1,528.22	18	307.37	11	1,501.21	16
	-	-	×	-	308.43	9	1,507.74	11	307.86	18	1,498.72	13
	-	-	-	×	306.67	5	1,496.02	6	307.00	5	1,495.93	9
	×	×	-	-	310.25	14	1,515.23	13	307.57	16	1,542.96	18
	×	-	×	-	309.18	11	1,496.80	7	307.67	17	1,498.66	12
	×	-	-	×	306.66	4	1,493.21	4	307.00	5	1,495.16	5
	-	×	×	-	310.40	15	1,519.58	15	307.50	13	1,498.87	14
	-	×	-	×	309.37	13	1,519.87	16	307.03	7	1,494.11	3
	-	-	×	×	306.57	3	1,492.17	3	307.27	10	1,495.43	7
	×	×	×	-	308.39	8	1,509.86	12	307.44	12	1,497.53	11
	×	×	-	×	309.26	12	1,518.92	14	306.97	4	1,496.35	10
	×	-	×	×	305.82	1	1,490.34	2	307.19	8	1,495.65	8
	-	×	×	×	308.99	10	1,494.23	5	307.21	9	1,495.24	6
×	×	×	×	306.20	2	1,486.36	1	306.94	3	1,493.77	2	
CGS	-	-	-	-	306.68	6	1,497.63	8	306.14	1	1,492.97	1
RAN	-	-	-	-	310.56	16	1,524.50	17	307.55	15	1,504.30	17
%GS	-	-	-	-	307.01	7	1,498.63	10	306.71	2	1,494.71	4

TABLE 2

Computational Results for the diverse and realistic dataset for 1,000 schedules

Diverse set										
	N25					N50				
	Avg_Pref	Avg_Pen	%Feas	Avg_CPU	%Dev_LP	Avg_Pref	Avg_Pen	%Feas	Avg_CPU	%Dev_LP
Case 1	241.51	53.02	88.27%	0.19	0.67%	481.73	84.81	89.95%	0.66	0.70%
Case 2	270.45	54.06	88.27%	0.69	0.33%	532.35	88.68	90.00%	2.00	0.38%
Case 3	251.31	53.02	87.86%	0.51	1.94%	500.91	84.98	89.03%	1.64	1.92%
Case 4	267.56	71.28	88.27%	0.77	0.51%	530.34	145.02	89.97%	2.39	0.56%
Case 5	242.56	53.02	85.82%	0.19	4.26%	484.01	84.83	84.77%	0.66	4.09%
Case 6	286.40	127.05	88.27%	1.73	0.77%	560.26	279.73	89.99%	4.56	0.87%
Case 7	260.88	72.78	79.68%	0.43	6.97%	515.59	142.59	77.39%	1.20	6.27%
Case 8	0.00	0.00	85.43%	0.00	3.99%	0.00	0.00	85.05%	0.00	3.66%
Realistic set										
	N75					N100				
	Avg_Pref	Avg_Pen	%Feas	Avg_CPU	%Dev_LP	Avg_Pref	Avg_Pen	%Feas	Avg_CPU	%Dev_LP
Case 1	739.70	150.34	88.67%	1.77	0.67%	1,183.72	166.57	90.30%	5.83	0.77%
Case 2	806.33	154.88	88.67%	3.56	0.35%	1,305.39	176.27	90.44%	8.19	0.45%
Case 3	754.12	150.40	87.89%	3.03	1.80%	1,213.13	167.63	89.14%	7.50	1.86%
Case 4	805.29	204.55	88.68%	4.14	0.48%	1,285.92	266.71	90.21%	8.72	0.61%
Case 5	741.23	150.36	85.84%	1.78	4.30%	1,186.81	166.68	85.58%	5.86	3.98%
Case 6	853.04	378.59	88.68%	7.06	0.37%	1,361.53	534.86	90.37%	12.33	0.48%
Case 7	790.97	211.52	78.66%	2.46	6.51%	1,263.28	264.58	78.52%	6.71	6.10%
Case 8	0.00	0.00	85.27%	0.00	4.83%	0.00	0.00	85.97%	0.00	4.65%
	N30					N60				
	Avg_Pref	Avg_Pen	%Feas	Avg_CPU	%Dev_LP	Avg_Pref	Avg_Pen	%Feas	Avg_CPU	%Dev_LP
Case 9	1,546.94	469.48	65.94%	4.71	-	3,156.39	801.67	68.23%	8.02	-
Case 10	1,437.84	396.77	69.17%	1.66	-	2,947.16	684.17	69.79%	6.84	-
Case 11	1,633.40	501.25	65.42%	13.33	-	3,321.11	846.15	68.13%	8.46	-
Case 12	1,476.42	399.48	69.17%	2.31	-	3,022.04	686.77	69.27%	6.87	-
Case 13	1,586.09	568.33	65.31%	5.11	-	3,237.34	987.08	67.40%	9.87	-
Case 14	1,453.73	402.81	68.75%	1.91	-	2,971.72	755.83	67.92%	7.56	-
Case 15	1,705.82	1,010.00	59.17%	18.04	-	3,483.89	2,085.52	59.58%	20.86	-
Case 16	1,556.91	545.42	65.73%	2.88	-	3,178.67	1,087.60	65.31%	10.88	-

TABLE 3

Computational Results for the diverse and realistic dataset for 5,000 schedules

Diverse set										
	N25					N50				
	Avg_Pref	Avg_Pen	%Feas	Avg_CPU	%Dev_LP	Avg_Pref	Avg_Pen	%Feas	Avg_CPU	%Dev_LP
Case 1	252.23	53.02	88.27%	1.61	0.19%	502.66	84.79	90.03%	5.33	0.25%
Case 2	240.86	53.02	88.27%	1.03	0.08%	480.42	84.79	90.03%	3.76	0.11%
Case 3	268.40	53.76	88.08%	2.31	1.11%	529.16	86.80	89.66%	6.38	1.17%
Case 4	250.33	53.02	88.27%	1.58	0.14%	499.12	84.79	90.03%	5.37	0.19%
Case 5	265.94	71.12	85.88%	2.26	3.63%	527.76	142.88	85.25%	6.97	3.44%
Case 6	241.86	53.02	88.27%	1.03	0.50%	482.65	84.79	90.03%	3.77	0.59%
Case 7	284.11	125.05	80.10%	5.05	5.95%	556.82	273.20	78.29%	12.61	5.33%
Case 8	259.21	71.89	85.60%	1.67	3.26%	513.02	140.01	85.50%	5.01	3.01%
Realistic set										
	N75					N100				
	Avg_Pref	Avg_Pen	%Feas	Avg_CPU	%Dev_LP	Avg_Pref	Avg_Pen	%Feas	Avg_CPU	%Dev_LP
Case 1	762.77	150.32	88.70%	10.13	0.33%	1,223.44	166.49	90.49%	23.08	0.36%
Case 2	738.20	150.32	88.70%	10.78	0.15%	1,180.94	166.32	90.51%	21.16	0.20%
Case 3	802.67	152.43	88.37%	14.33	1.24%	1,299.15	170.40	90.01%	24.68	1.22%
Case 4	752.10	150.32	88.70%	9.83	0.23%	1,209.59	166.45	90.48%	22.09	0.28%
Case 5	802.08	202.88	86.06%	11.36	3.85%	1,280.83	260.23	86.28%	24.08	3.43%
Case 6	739.66	150.32	88.70%	8.51	0.17%	1,183.93	166.35	90.51%	21.22	0.22%
Case 7	848.60	367.17	79.48%	16.28	5.76%	1,354.75	517.49	79.49%	31.24	5.38%
Case 8	787.75	206.68	85.78%	9.42	4.30%	1,258.08	256.98	86.53%	22.11	4.11%
Realistic set										
	N30					N60				
	Avg_Pref	Avg_Pen	%Feas	Avg_CPU	%Dev_LP	Avg_Pref	Avg_Pen	%Feas	Avg_CPU	%Dev_LP
Case 9	1,525.63	423.65	68.02%	17.37	-	3125.66	743.65	68.54%	37.99	-
Case 10	1,429.45	392.50	69.58%	7.52	-	2935.91	677.81	70.00%	24.01	-
Case 11	1,607.99	429.58	67.60%	40.68	-	3277.91	758.54	68.44%	73.00	-
Case 12	1,465.36	392.40	69.38%	8.95	-	3006.31	677.29	69.90%	26.10	-
Case 13	1,570.33	501.67	66.35%	17.52	-	3211.65	906.46	67.92%	37.66	-
Case 14	1,444.48	394.17	69.48%	7.64	-	2963.95	681.98	69.79%	24.04	-
Case 15	1,678.09	843.75	63.33%	59.68	-	3424.43	1543.33	66.04%	105.97	-
Case 16	1,541.34	498.02	66.88%	10.41	-	3154.49	882.92	68.23%	27.44	-