Real-Time Scheduling for Nurses in Response to Demand Fluctuations and Personnel Shortages

Jonathan F. Bard[1], Hadi W. Purnomo[1]

[1]Graduate Program in Operations Research
& Industrial Engineering
The University of Texas
ETC 5.160, C2200
Austin, Texas 78712, USA
jbard@mail.utexas.edu, hadiwaskito@yahoo.com

Abstract. This paper discusses the problem that hospital managers face several times a day in response to the demand for nursing services. Prior to the start of each shift, the number of nurses who are scheduled to be on duty over the next 24 hours is compared with the number actually available, and if shortages exist a series of decisions have to be made to ensure that each unit in the hospital has sufficient coverage. These decisions involve the use of overtime, outside nurses, and floaters. To address this problem, we have developed an integer programming model that takes the current rosters for regular and pool nurses and the expected demand for the upcoming 24 hours as input, and produces a revised schedule that makes the most efficient use of the available resources. Problem instances with up to 120 nurses are shown to be solvable in a negligible amount of time.

1 Introduction

Most of the recent research on nurse scheduling has concentrated on rostering with the aim of accommodating individual preferences. Examples of preferences include requests to work specific shifts or to be given specific days off, and can be measured in terms of the number of working hours, shift sequence patterns or even nurse to patient ratios (see [6] for a survey). As a first step, nurses are typically asked to sign up for shifts prior to the beginning of the planning horizon. At that time, they may also submit a list of requests to the nurse manager who decides which to approve immediately and which to defer in light of expected demand. The outcome is a midterm schedule for each nurse in the hospital.

Midterm scheduling fixes the work assignments for the permanent nursing staff for up to six weeks at a time. Each unit generates its own rosters independently using some measure of “average” demand as input. Most approaches are based on integer programming formulations coupled with heuristics; see, e.g. [1], [2], [3], [7], [9]. Petrovic et al. [8]; however, use an artificial intelligence methodology called case-based reasoning that aims to imitate human style decision making by solving new problems using knowledge about solutions to similar problems.

In this paper, we begin with the midterm schedule and address the problem of adjusting individual work assignments to account for daily fluctuations in the patient population, absenteeism, and emergencies. Possible options include the use of overtime, calling in nurses on their day off, using outside resources and pool nurses, or living with the shortages.

The problem is formulated as an integer program and solved within a rolling horizon framework that spans 24 hours. Solutions are obtained prior to the beginning of each 8-hour shift. Testing indicated that instances with up to 120 nurses can be solved in a few second in most cases. This is extremely important in an operational environment where the situation changes continually.

In the next section, we describe the daily adjustment problem and discuss the roles and responsibilities of the nurse management team. The presentation is based on our experience at several medium size hospitals in the U.S. but provides as much generalization as possible. In Section 3 we present the mathematical model for the problem. The data requirements are discussed in Section 4, and computational results for instances with up to 120 nurses working in 14 units are highlighted in Section 5. We close with some remarks on implementation and the effectiveness of the approach.

2 Problem Statement

The dynamic nature of the demand for nursing services, coupled with sick leave, personal days, and emergencies, requires the midterm schedule to be adjusted on a shift-by-shift basis. In most hospital, this rescheduling is performed throughout the day a few hours prior to the start of each standard 8-hour shift – day (D = 7:00 a.m. – 3:00 p.m.), evening (E = 3:00 p.m. – 11:00 p.m.), night (N = 11:00 p.m. – 7:00 a.m.). To enable the process, patient acuity and census data are periodically entered into the hospital's management information system. The primary information includes admissions, discharges and transfers (ADT), and projected ADTs for the next shift.

In forecasting demand, a major difficulty lies not only in determining the number of patients for whom care must be provided, but the level of care that will be required by each; i.e., the acuity of care. Siferd and Benton [9] represent the number of nurses needed during a shift as a multiplicative model of mean patient acuity, number of patients, and the mean rate of change in acuities. Using simulation, they show that the stochastic interplay of these factors can to cause wide swings in coverage requirements for subsequent shifts. For issues associated with staff scheduling in a retail environment where demand varies stochastically over the day, see, e.g., [4].

To make the staffing decisions as the day unfolds, it is necessary to have a "24-hour view" for each unit. The additional information required is the schedule for the next 24 hours, the list of call outs, the list of expected ADTs, the list of available nurses (on call, pre-approved part-timers, casuals, floaters, agency options), and the projected demand for nurses by skill type. By way of clarification, floaters are nurses who are generally employed by the hospital but not assigned to a specific unit; however, they are given a monthly schedule that specifies the shifts that they will be working. When they arrive at the hospital on their assigned days, they are told where to report. In contrast, casuals have no employment contract and so are not given a monthly schedule. They make their availability known to either unit managers or the nursing services office, and are typically called at the last minute.

In general, the nursing resources director in the hospital is responsible for ensuring that all units are covered. At his or her control are the float pool, casuals and other external resources. At the unit level, the nurse manager or clinical manager has the authority to assign overtime, request adjustments to current schedules, and call in nurses not working that day. It is either not permitted or extremely undesirable to alter a nurse's midterm schedule without his or her consent. To a large extent, this restriction limits the potential cost savings in the short run but promotes normalcy and stability in the long run.

The goal of the rescheduling effort is to reallocate the available resources in a way that minimizes the cost of the shortfall. In doing so, it is important to minimize the differences between the new plan and the original plan. In the hospital environment, this often leads to conflict because the optimal course of action may impose undesirable schedules on the permanent staff, such as excessive overtime and long work stretches.

2.1 Decision Makers

The specific people involved in decision making, along with their roles, depends on the particular option under consideration, predefined authorizations, coverage requirements, and time of day. At the unit level, the nurse manager, clinical manager, or nurse in charge keeps track of the current situation and assesses whether the level of coverage is under, appropriate or over for the number of in-patients and their acuity. This information is reported to the supervisors who among other things, call on the float pool when necessary to fill in when shortages exist. To facilitate their role, supervisors are provided with worksheets that are updated by shift and show who is scheduled for duty in each of the units in their clinical areas.

The nursing resources director in conjunction with the nursing services office manages the external resources that are available to the hospital. These include what we will collectively call *outside nurses* comprising agency nurses, casuals, other per diem, and perhaps an external float pool. The monthly schedules for the outside nurses and internal float pool nurses are readily available to nursing services. The latter are provided by the float pool managers for critical care and med/surg. Specific unit assignments are not made at this point, but are left to the supervisors as daily demand dictates. Most hospitals do not have float pools for units outside of these areas.

2.2 Daily Adjustments

When more nursing power is available in a unit than is needed for a particular shift, the nurse manager or equivalent has several options. Each is exercised in turn beginning with the least senior staff member. The first option is to try to float the nurse to another unit or to reassign her to a later day in the same pay period (usually 14-day blocks). If the nurse is not willing to float or be reassigned, the shift is cancelled. In either case, the supervisor is notified of the situation. In critical care units at the hospitals that we are working with, if a nurse is cancelled, she is placed on the on-call list for some period of time (med/surg and the other units do not use this rule). Generally speaking, nurses would rather be cancelled than floated to another unit. If the hospital is requesting the cancellation, then one of the following designations is used for the time off: vacation, personal day, holiday, or unpaid leave. Each has different cost consequences.

When shortages exist in a unit, a number of steps can be taken to compensate, each also having different cost consequences. In most hospitals, though, cost is not foremost on the mind of the person directly responsible for the unit. The primary objective is to achieve sufficient coverage, especially in situations that are critical and in which coverage requirements are mandated by law. With this in mind, the order of action is typically:

1. Look for a volunteer in the unit to work the next shift (or fraction thereof) as overtime.
2. Try to reach casuals or per diem nurses.
3. Have the supervisor either negotiate for floaters with other units that might be overstaffed or draw on the float pool.
4. Try to reach unit staff who are not scheduled to work during the current day.
5. Cycle through the on-call list (used mostly in emergencies)
6. Have the nursing resources director call in agency nurses.
7. Invoke mandatory overtime by requesting that a nurse on the current shift stay for the next shift.

Although it is helpful to have guidelines, this seven-step procedure can rarely be followed exactly without some judgment being exercised. For example, when a nurse is called in on her day off to work a shift that is separated by her upcoming shift by only 8 hours (e.g., evening → day), she is likely to call in sick for the day shift so little has been gained.

3. Mathematical Model

The majority of rules and constraints that govern daily scheduling have been mentioned above. The complete set may vary among hospitals, but generally reflects institutional policies, union agreements, state or federal statutes, and financial considerations. The over goal is to satisfy coverage requirements at minimum cost while taking into account nurse preferences, morale, the need for the perception of fairness, and the expected response of staff members whose work patterns are affected. Bearing this in mind, the problem will be formulated as an integer linear program for a predetermined planning horizon of, say, 24 hours or 3 shifts.

In our approach, solutions will be obtained using a rolling horizon strategy; i.e., the problem is solved for the 3 upcoming shifts (say, D, E, N) and the results implemented for at least the next shift (D) and perhaps for all 3 shifts (E, N as well). When the current shift (D) expires and the next shift (E) begins, the problem is re-solved for the next 3 shifts (E, N, D), and so on. It is a simple matter to include 12-hour shifts in the model as long as their start times coincide with one of the 8-hour shifts. We denote them by AM (typically 7:00 a.m. – 7:00 p.m.) and PM (7:00 p.m. – 7:00 a.m.).

In formulating the model, it is assumed that all costs associated with assigning a nurse to work overtime or an off day are known, that demand is given or can be accurately estimated for each of the three shifts in the planning horizon, and that the status of all unit nurses, pool nurses, casuals, and agency nurses is known. This means that the nurse managers, the supervisors, and the nursing resources director all have up-to-date information on call outs, shortages, surpluses, floaters, and pool nurses. Although not critical, it is also assumed that the problem decomposes by subsets of units, such as critical care and med/surg, as well as by skill type.

As we will see, the model reflects the point of view of the hospital and is intended for use by the nursing services office rather than the unit managers. Prior to running the model, the unit managers are expected to evaluate their current staffing needs and take some action if necessary. If a staffing shortage in view, they will either ask one or more of their nurses currently on the floor to work overtime, or try to reach casual or per diem nurses which whom they have a last minute arrangement. These decisions are not included in the model; however, if a nurse is not needed in her home unit and is willing to float before, during, or after a regularly scheduled shift, this option is included.

3.1 Notation

The following notation is used in the developments.

Indices

- i index for nurses
- j, k index for units
- s index for shifts
- p index for periods

Sets

- J set of units under consideration
- T set of time periods in planning horizon
- $J(i)$ set of units in which nurse i is qualified to work (outside nurses include casuals, agency nurses, and other)
- $J(i, j)$ set of units to which nurse i currently working in unit j can float
- S set of shifts in planning horizon
- $S(i)$ set of shifts nurse i is permitted to work other than the shift she is assigned in the midterm schedule; could include 4-hour overtime sifts
- $\hat{S}(i)$ shift(s) nurse i is assigned to work in the midterm schedule; $\hat{S}(i) \subseteq S \setminus S(i)$
- R set of regular nurses
- $R(j)$ set of regular nurses that can work in unit j

P	set of pool nurses
$P(j)$	set of pool nurses that can work in unit j

Input data

c_{ijs}	cost of assigning nurse i to unit j for shift s (may be written as c_{ijs}^{θ} on occasion, where θ designates the type of cost; e.g., agency, pool nurse, day off, float)
c_{ijks}	cost of floating nurse i from unit j to unit k for shift s
c_{jp}	cost of assigning on-call nurse to unit j during period p
D_{jp}	incremental number of nurses required in unit j for period p (+ means shortage, – means surplus)
a_{ps}	parameter equal to 1 when shift s covers period p; 0 otherwise
M	large penalty coefficient
c_i^1	cost of unproductive assignment (cancellation) for nurse i
c_i^2	incremental cost for floating nurse i and then assigning overtime on the next shift
p_i^1	penalty for floating nurse i to another unit
p_i^2	penalty for floating nurse i and then assigning overtime on the next shift
p_i^3	penalty for unproductive assignment (cancellation) of nurse i
p^4	penalty for on-call assignment
$P^{\max}$	maximum number of total undesirable patterns allowed

Decision variables

x_{ijs}	1 if nurse i in unit j is assigned overtime during shift s; 0 otherwise
y_{ijs}	1 if pool nurse i is assigned to unit j during shift s; 0 otherwise
z_{js}	number of outside nurses assigned to unit j during shift s
o_{jp}	number of on-call nurses assigned to unit j in period p
w_{ijks}	1 if nurse i who is assigned to unit j in the midterm schedule floats to unit k during shift s; 0 otherwise
u_i	1 when nurse i is floated and then assigned overtime on two consecutive shifts; 0 otherwise
v_i	1 if nurse i (regular or pool) is not needed on the shift assigned to him or her in the midterm schedule; i.e., the shift is canceled
g_{jp}	number of gaps (uncovered demand) in unit j in period p

3.2 Formulation

The 0-1 integer programming model for a fixed planning horizon and single skill type is as follows.

Minimize $\sum_{i\in R}\sum_{j\in J(i)}\sum_{s\in S(i)} c_{ijs}x_{ijs} + \sum_{i\in P}\sum_{j\in J(i)}\sum_{s\in S(i)} c_{ijs}y_{ijs}$

$$+\sum_{j\in J}\sum_{s\in S} c_{js}z_{js} + \sum_{i\in R}\sum_{j\in J}\sum_{k\in J(i,j)}\sum_{s\in \hat{S}(i)} c_{ijks}w_{ijks} + \sum_{i\in R\cup P} c_i^1 v_i$$

$$+\sum_{i\in R\cup P} c_i^2 u_i + \sum_{j\in J}\sum_{p\in T} c_{jp}o_{jp} + M\sum_{j\in J}\sum_{p\in T} g_{jp} \quad (1a)$$

subject to $\sum_{i\in R(j)}\sum_{s\in S(i)} a_{ps}x_{ijs} + \sum_{i\in P(j)}\sum_{s\in S(i)\cup\hat{S}(i)} a_{ps}y_{ijs}$

$$+\sum_{i\in R}\sum_{j\in J}\sum_{k\in J(i,j)}\sum_{s\in\hat{S}_i} a_{ps}w_{ijks} + \sum_{j\in J}\sum_{s\in S} a_{ps}z_{js}$$

$$+ o_{jp} + g_{jp} \geq D_{jp}, \forall j\in J, \forall p\in T \quad (1b)$$

$$\sum_{j\in J(i)}\sum_{s\in S(i)} x_{ijs} \leq 1, \forall i\in R \quad (1c)$$

$$\sum_{j\in J}\sum_{k\in J(i,j)}\sum_{s\in\hat{S}(i)} w_{ijks} + v_i = 1, \forall i\in R \quad (1d)$$

$$\sum_{j\in J(i)}\sum_{s\in S(i)} x_{ijs} + \sum_{j\in J}\sum_{k\in J(i,j)}\sum_{s\in\hat{S}(i)} w_{ijks} - u_i \leq 1, \forall i\in R \quad (1e)$$

$$\sum_{j\in J(i)}\sum_{s\in\hat{S}(i)} y_{ijs} + v_i = 1, \forall i\in P \quad (1f)$$

$$\sum_{j\in J(i)}\sum_{s\in S(i)} y_{ijs} \leq 1, \forall i\in P \quad (1g)$$

$$\sum_{i\in R}\sum_{j\in J}\sum_{j\in J(i,j)}\sum_{s\in S(i)} p_i^1 w_{ijks} + \sum_{i\in R} p_i^2 u_i + \sum_{i\in R\cup P} p_i^3 v_i$$

$$+\sum_{j\in J}\sum_{p\in T} p^4 o_{jp} \leq P^{\max} \quad (1h)$$

$x_{ijs}, y_{ijs} \in \{0, 1\} \forall i, j, s, , w_{ijks} \in \{0, 1\} \forall i, j, k, s,$

$z_{js} \geq 0$ and integer $\forall j, s, o_{jp}, g_{jp} \geq 0$ and integer

$\forall j, p, u_i, v_i \in \{0, 1\} \forall i$ (1i)

The objective function (1a) sums the costs of each alternative available for handling shortages. The first term is associated with the permanent staff and covers the possibility of voluntary overtime and calling a nurse in on her day off if she is not already on the on-call list. These options are mutually exclusive enforced by constraint (1c) so the decision to assign nurse i to shift s in unit j is unique. Referring to the cost coefficients, if it is more expensive to call nurse i in on her day off than to assign her an overtime shift, we would require $c_{ijs}^{\text{OFF}} > c_{ijs}^{\text{OT}}$, although the values used in the model do not have to be the actual costs.

The second and third terms in (1a) account for the use of nurses from the internal float pool and from outside sources, respectively. Depending on availability and cost, either or both options can be part of a solution. The fourth term considers all nurses scheduled to work shift s in unit j and allows those that are willing and qualified to float to unit $k \in J(i,j)$. The set $J(i,j)$ must be defined dynamically to ensure that if nurse i is scheduled to work shift s and is needed in her home unit, she is not floated even if she is eligible. The fifth and sixth terms respectively account for the cost of canceling a nurse and the extra payment associated with successively being floated and then working overtime. The seventh term is the cost for using on-call nurses, the last term is associated with the gaps.

Constraint (1b) ensures that demand is met in every unit $j \in J$ for each $p \in T$ in the planning horizon. Instead of using shifts as the unit time, the constraint is written in terms of periods. This representation is needed whenever two shift-types overlap, which is the case when 8-hour and 12-hour shifts are included in the model. When $D_{jp} = 0$, no action is necessary; when $D_{jp} < 0$, there is over coverage in unit j. This is addressed in constraint (1d). When $D_{jp} > 0$, the under coverage can be made up by various options. Because nurses work in terms of shifts and shifts cover several periods, the decision to float a nurse must be based on the coverage situation in each of those periods. In particular, extra resources must be available throughout the shift if the floating option is to be considered.

Constraint (1c) ensures that each regular nurse i is given at most one overtime assignment during the planning horizon. That assignment must be for a single unit. Constraint (1d) ensures that a nurse can only be floated to one unit on a shift. Recall that the decision as to which nurse to float is based on a rotating list. Thus, seniority becomes less important as time passes. The option to float or to call in a nurse on her day off for one shift and then assign her overtime on the next shift is permitted by the model. Specifying the data correctly is necessary to preclude the possibility of non-contiguous assignments.

When a nurse is not needed for her scheduled shift and there is no unit to which she can be floated, she is said to be *unproductive*. Depending on hospital rules and contractual agreements, it may be possible to cancel a nurse, usually at some cost and with some penalty. The variable v_i affords this option. The actual cost depends in part on whether a pool nurse or a regular nurse is under consideration. The variable u_i in constraint (1e) indicates when nurse i floats during one shift and then works overtime on another. This situation incurs an extra cost c_i^2, as indicated by the fifth objective function term. It is also undesirable from the nurse's point of view and so is included as a preference violation term in constraint (1h).

Constraint (1f) limits the assignment of each pool nurse $i \in P$ to no more than one unit during her midterm schedule shift. Like constraint (1d), the cancellation variable v_i is also included in the constraint to capture the situation in which the nurse is not assigned to any unit. Constraint (1g) limits the number of overtime shifts that a pool nurse can work to one. When this option is not available, the set $S(i)$ is empty and the constraint is omitted from the model.

Constraint (1h) is designed to take into account preference violations in the adjusted schedule. The intent is to restrict the total number of undesirable patterns to no more than a user-supplied parameter, $P^{\max}$, as well as limit the use of on-call nurses (fourth term on left). Three different types of undesirable patterns are considered in the model: (1) floating a regular nurse from her home unit during her scheduled shift, (2) canceling a nurse, and (3) floating a regular nurse and then assigning overtime for the next shift.

Finally, constraint (1i) defines the domain of the decision variables. In practice, upper bounds exist on the number of on-call nurses o_{jp} available in period p and the number of outside nurses z_{js} available for shift s in each unit j. As an aside, we note that it is not really necessary to distinguish between the regular and pool nurse variables, x_{ijs} and y_{ijs}, because the sets $S(i)$, $J(i, j)$ and $P(j)$ are uniquely defined. Doing so, however, makes the model easier to understand.

4. Parameter Settings and Data Requirements

In order for the solution of model (1a) – (1i) to mimic the sequential decision-making process outlined in Section 3, the cost coefficients in (1a) must be defined appropriately. For example, if pool nurses are to be used before voluntary overtime, then $c_{i_1 js}^{\text{Pool}} < c_{i_2 js}^{\text{OT}_\text{V}}$ for all $i_1 \in P(j), j \in J(i_1), s \in S(i_1)$ and $i_2 \in R, j \in J(i_2), s \in S(i_2)$. Similarly, if it is desirable to use nurses who are off before assigning mandatory overtime, then $c_{ijs}^{\text{OFF}} < c_{ijs}^{\text{OT}_\text{M}}$ for all $i \in R \cup P, j \in J(i), s \in S(i)$. These coefficients may reflect actually costs or may be set artificially to enforce a predetermined selection order. In any case, when reporting the final solution, the true values should be used in the calculations.

The basic time unit in the model is 4 hours corresponding to the largest time increment that evenly divides 8- and 12-hour shifts. On-call nurses are already assigned in 4-hour blocks as are gaps and overtime, but if it were desirable to assign outside hours in increments of say, 2- or 4-hour blocks, then the regular, pool nurse and float variables would have to be defined accordingly. This would greatly increases the size and complexity of the model, as would allowing for split shifts. Although the incremental demand data, D_{jp}, is currently specified in 4-hour blocks, it would have to be further disaggregated to account for, say, 2-hour overtime assignments or shift starting times other than those that coincide with the basic 6 periods.

Critical to the successful application of model (1a) – (1i) to the daily scheduling problem is (near) real-time, automated data input and updating. As the day progresses, the sets used in the formulation of the model change, as do the demand and cost coefficients. If the burden is too great at the beginning of each shift to input data, the nurse manager or designee who will be running the system is likely to abandon it in favor of the current manual procedures.

Fortunately, many of the lists and sets needed to run the model are static so, at most, only minor updates will be required between runs. For example, the list of outside nurses who are available by unit is known on a monthly basis, so it would only be necessary to plug in those associated with the current planning horizon. Most of the other input data are either static or readily available from the human resources database, perhaps with the exception of the voluntary overtime list. One way to mitigate this problem is to ask the nurses during the midterm sign-up period to indicate which shifts they would be willing to work overtime. As the month unfolds and overtime is accumulated, the voluntary overtime list would be updated by either the nurse manager or the nurses themselves as the situation changed.

5. Computational Results

The model was tested by solving a range of problems for a 14-unit hospital with a staff of approximately 300 regular and pool nurses. Depending on the instance, between 40 and 120 nurses were candidates for rescheduling over the 24-hour planning horizon. All codes were written in C++ language and run on a PC with a Pentium 1.3 GHz processor. The IPs were solved with the CPLEX callable libraries.

5.1 Input Data

The data in Table 1 summarizes the seven problem instances investigated. In each case, we set the number of pool nurses to 20 and varied the number of regular nurses. The second column indicates the total number of nurses in each problem; for example, problem 1 contains 20 regular nurses and 20 pool nurses. The regular nurses that are included as part of the input are only those available to work overtime or to be floated to other units in the next 24-hours. The third and fourth columns indicate the number of constraints and variables in the corresponding IPs. Most of the variables are binary representing the assignment of a nurse to a particular unit for a particular shift. We do not allow split-shift assignments, where a nurse works for two different units within her 8- or 12-hour shift. The next three columns indicate the total staff shortfall, the number of agency nurse available, and the number of nurses on call, respectively. All values are given in terms of 8-hour shifts. For problem 1, for example, in the upcoming 24 hours, 69 shifts are uncovered, 18 agency nurses are available, and 26.5 nurses are on call. The last column indicates the average number of units to which a regular nurse can float. Pool nurse are eligible to work in approximately 7 different units.

Table 1. Properties of test data

Problem No.	Total no. nurses	No. of con-traints	No. of vari-ables	Staff short-fall	No. agency nurses	On-call nurses	Average units to float
1	40	185	634	69	18	26.5	2.85
2	70	275	1024	92	12	8	3.21
3	90	285	1120	92	13	16	3.33
4	40	185	671	53.5	9	7.5	5
5	70	270	1127	76	12	8	3.72
6	90	285	1328	92	13	16	3.85
7	120	425	1871	100	13	16	5

In defining a problem instance, we begin by specifying a credential unit list for each nurse that indicates the units in which she is eligible to work. The input data for pool nurses and regular nurses is given in Tables 2 and 3, respectively, for problem 1. The second column in Table 2 identifies the current shift assignment for each pool nurse. There are 5 possibilities: AM, PM, Day, Evening and Night. The unit credential list is given for each nurse in the next column, followed by her hourly wage in US dollars. These values are used to determine the overtime cost coefficients, c_{ijs}, and the float cost coefficients, c_{ijks}, in (1a).

The input in Table 3 is similar to that of Table 2, with some additional parameters. The home unit for each nurse is given in column three and the overtime periods she is allowed to work are given in column five. For the latter, the shift designations are slightly different due to the existence of 4-hour overtime. The qualifiers "Early" and "Late" are used to distinguish this case.

Table 4 lists the value of each parameter in constraint (1h) and the cost coefficients used in the objective function (1a). The entries for c_{jp} and c_{js} are constant for all indices implying that there is no differentiation among on-call nurses and among agency nurses. For the last entry, c_{ijks}, only the base value of is given. The actual value used was obtained by perturbing the hourly rate by a small amount to differentiate the cost of reassigning nurse i to any of the units in which she is eligible to work.

Table 2. Input data for pool nurses for problem 1

Pool nurse, i	Assigned shift, $\hat{S}(i)$	Allowed units, $S(i)$	Hourly wages ($)
1	AM	0, 1, 2, 3, 4, 5, 6	17
2	PM	1, 2, 3, 4, 5, 6, 7	17
3	Day	2, 3, 4, 5, 6, 7, 8	18.5
4	Evening	3, 4, 5, 6, 7, 8, 9	16
5	Evening	4, 5, 6, 7, 8, 9, 10	17
6	Night	5, 6, 7, 8, 9, 10, 11	19
7	Night	6, 7, 8, 9, 10, 11, 12	18.5
8	Day	7, 8, 9, 10, 11, 12, 13	18.5
9	Day	8, 9, 10, 11, 12, 13, 0	19
10	Day	9, 10, 11, 12, 13, 0, 1	18
11	Evening	10, 11, 12, 13, 0, 1, 2	17
12	Evening	11, 12, 13, 0, 1, 2, 3	17
13	Night	12, 13, 0, 1, 2, 3, 4	17
14	Night	13, 0, 1, 2, 3, 4, 5	17
15	Day	0, 1, 2, 3, 4, 5, 6	18.5
16	Day	1, 2, 3, 4, 5, 6, 7	18
17	Evening	2, 3, 4, 5, 6, 7, 8	19
18	Evening	3, 4, 5, 6, 7, 8, 9	17
19	Night	4, 5, 6, 7, 8, 9, 10	17
20	Night	5, 6, 7, 8, 9, 10, 11	17

Table 5 gives the availability data for agency nurses by shift and for on-call nurses by period for problem 1. The use of the latter is restricted to evening and night shifts. Table 6 displays the incremental demand data, D_{jp}, by unit and period, where a value of 0 actually means that $D_{jp} \leq 0$. When a surplus exists, it is assumed that the nurse manager has identified those individuals who will be either floated or cancelled in the solution.

Table 3. Input data for regular nurse for problem 1

Regular nurse, i	Assigned shift, $\hat{S}(i)$	Home unit	Allowed units, $S(i)$	Hourly wage ($)	Permitted over-time periods
1	Off	1	0, 1, 3, 4, 5	17	Day
2	Day	2	1, 3, 6, 7	18.5	Evening
3	Day	4	3, 5, 6, 7	21	Evening
4	Evening	6	2, 10, 12	17	Day, Night
5	Evening	11	3, 11, 12	18.5	Day, Night
6	Night	2	10, 11, 12	22	Evening
7	Night	2	7, 8, 9	21	Evening
8	Off	1	0, 6, 7	21	Evening
9	Day	2	4, 8, 9	18.5	Evening
10	Day	6	5.9, 10	17	Evening
11	Evening	10	8, 9, 10	17	Day, Night
12	Evening	10	8, 9, 12	18.5	Day, Night
13	Night	9	8, 9	18.5	Late evening
14	Night	9	9, 10	18.5	Late evening
15	Day	9	0, 4, 5	21	Early evening
16	Day	9	0, 4, 7	21	Early evening
17	Evening	6	6, 7, 8, 9	22	Late day
18	Evening	13	6, 7, 8, 9	17	Early night
19	Night	9	10, 11	17	Evening
20	Night	2	10, 11, 13	21	Evening

Table 4. Model parameters and their values

Parameters	Value
p_i^1 penalty for floating nurse i to another unit	1
p_i^2 penalty for floating nurse i and then assigning overtime in next shift	4
p_i^3 penalty for unproductive assignment (cancellation) of nurse i	6
p^4 penalty for an on-call assignment	8
c_i^1 cancellation cost for nurse i	3 hr of pay
c_i^2 incremental cost of floating nurse i and then assigning over-time in next shift	$10
c_{jp} cost of using an on-call nurse in unit j during period p	$120/period
c_{js} cost of using an agency nurse in unit j during shift s	$312
c_{ijks} reassignment cost from unit j to unit k for shift s for nurse i	hourly rate

Table 5. Availability data for agency and on-call nurses for problem 1

	Agency shift			On-call period					
Unit	Day	Evening	Nigh t	1	2	3	4	5	6
0	1	1	1	0	0	0	1	2	1
1	1	1	1	0	0	0	1	1	2
2	0	0	0	0	0	0	1	1	1
3	0	0	0	0	0	0	2	1	1
4	1	1	1	0	0	0	1	2	1
5	1	1	1	0	0	0	1	1	2
6	0	0	0	0	0	0	1	1	1
7	0	0	0	0	0	0	2	1	1
8	1	1	1	0	0	0	2	1	1
9	1	1	1	0	0	0	1	1	2
10	0	0	0	0	0	0	1	1	1
11	0	0	0	0	0	0	2	1	1
12	0	0	0	0	0	0	2	1	1
13	0	0	0	0	0	0	1	2	1

Table 6. Incremental demand by unit for problem 1

	Period					
Unit	1	2	3	4	5	6
0	2	2	1	1	2	2
1	1	1	2	2	2	2
2	0	0	3	3	0	0
3	1	1	2	2	1	1
4	0	0	0	0	0	0
5	1	1	3	3	4	4
6	4	4	0	0	4	4
7	2	2	2	2	1	1
8	3	3	2	2	3	3
9	0	0	3	3	0	0
10	3	3	0	0	3	3
11	2	2	0	0	2	2
12	2	2	2	2	2	2
13	1	1	0	0	2	2

5.2 Experimental Design and Results

Two sets of experiments were run to determine the computational effort required to solve model (1a) – (1i) and to investigate the implications of trying to accommodate preferences on a daily basis. In the first set, the preference constraint (1h) is omitted so cost is the only consideration. In the second set, we studied the tradeoff between monetary outcome and preference violations by parametrically varying the value of P^{max} in (1h). Initially, P^{max} is set to the value associated with the solution obtained in the first set of experiments in which (1h) is not included (see last row in Table 7).

Computation times were negligible for all instances. Solutions were found by CPLEX at the root node of the search tree within a fraction of a second for the first set of experiments and within several seconds for the second set. At first, we believed that this was primarily due to the lack of differentiation in the cost c_{ijs} of assigning nurse i to unit j for all $j \in J(i)$ during shift s, as well as the lack of differentiation in the float cost c_{ijks} as originally defined. Both sets of coefficients are a function of the wage rate given in Table 2 for nurse i and initially were assumed to be independent of unit assignments j and k. Perturbing the values of c_{ijs} and c_{ijks}, though, had no effect on the computational effort or the size of the search tree.

One possible explanation for this relates to the structure of the local constraints associated with each nurse i. Although the feasible region of the IP is not totally unimodular, Eqs. (1c) – (1g) can be rewritten as flow balance constraints by appropriately redefining the decision variables. Anecdotally speaking, the presences of a pure network substructure in a problem often leads to quick solutions. A summary of the results is given in Table 7 for the first set of experiments. The gaps are filled initially by the least expensive option, the pool nurses, and then by combinations of the remaining resources. In the absence of a restriction on preference violations, sufficient resources are available to satisfy all demand, expect for problem 7. Solutions were always found at the root node, after several dozen rows and columns were eliminated by CPLEX's presolve routine, and after an equal number of constraints were added by the cut generator. These cuts were derived from the intersection graph constructed from a portion of the model's A-matrix.

Table 7 also reports the quality of the solution for each problem instance as measured by the number of pool nurses used, the number of regular nurses floated, the number of overtime 8-hour shifts included, and the extend to which on-call and agency nurses are used. The bottom row of the table gives the weighted sum of undesirable patterns associated with each solution. We designate this value P^{max}.

When the preference constraint is activated by reducing the value of P^{max} to some number below the maximum determined in the first set of experiments, the feasible region becomes tighter, thus restricting the set of feasible solutions. Computations times increase a bit because more cuts in the form of cover inequalities are added by CPLEX in the enumeration process before convergence occurs. Nevertheless, total computational times were no more than a few seconds and only a handful of nodes had to be explored.

The assignments resulting from solution of problem 1 without constraint (1h) are shown in Tables 8 through 11. Table 8 reports the unit assignment for each pool nurse as well as the input data previously given in Table 2. This redundancy is for ease of comparison. Table 9 describes the assignments for the regular nurses. The first four columns are part of the input, the fifth and sixth columns indicate the overtime assignments in terms of the unit and period, respectively. The floating assignments are shown in the last column of Table 9.

Table 7. Summary of computations for first set of experiments

Output features	Problem no.						
	1	2	3	4	5	6	7
Pool nurses used	20	20	20	20	20	20	20
Regular nurses floated	17	37	44	17	38	48	66
Overtime nurses (shifts)	7	5	11	1	2	5	11
On-call nurses used (shifts)	9	4	4	3	5	8	1
Agency nurses needed (shifts)	13	10	12	8	11	10	6
Gaps	2.5	0	0	0	0	0	1
No. of nodes in B&B tree	1	1	1	1	1	1	1
Solution time (sec)	< 1	< 1	< 1	< 1	< 1	< 1	< 1
Cost	$17,172	$11,869	$16,152	$8,519	$12,231	$15,426	$19,705
P^{max} at solution	258	308	442	138	310	424	500

As an example, consider nurse 4 who is was assigned to work an Evening shift in the midterm schedule. Instead she is floated from her home unit 6 to unit 2 for the Evening shift. Also, because she was available for overtime and several of the units in which she is eligible to work were understaffed, the solution assigned her a Day shift in unit 10. This scenario illustrates one of the more complicated situations.

The assignments for the on-call and agency nurses are provided in Table 10. It is assumed that each agency nurse must be hired for an 8-hour shift. The last six columns in the table give the assignments for the on-call nurses for each 4-hour period in the planning horizon.

Table 8. Example assignment for pool nurses in problem 1

Pool nurse	Assigned shift	Allowed unit	Assigned unit
1	AM	0, 1, 2, 3, 4, 5, 6	3
2	PM	1, 2, 3, 4, 5, 6, 7	3
3	Day	2, 3, 4, 5, 6, 7, 8	8
4	Evening	3, 4, 5, 6, 7, 8, 9	7
5	Evening	4, 5, 6, 7, 8, 9, 10	7
6	Night	5, 6, 7, 8, 9, 10, 11	6
7	Night	6, 7, 8, 9, 10, 11, 12	6
8	Day	7, 8, 9, 10, 11, 12, 13	13
9	Day	8, 9, 10, 11, 12, 13, 0	12
10	Day	9, 10, 11, 12, 13, 0, 1	11
11	Evening	10, 11, 12, 13, 0, 1, 2	1
12	Evening	11, 12, 13, 0, 1, 2, 3	12
13	Night	12, 13, 0, 1, 2, 3, 4	1
14	Night	13, 0, 1, 2, 3, 4, 5	5
15	Day	0, 1, 2, 3, 4, 5, 6	6
16	Day	1, 2, 3, 4, 5, 6, 7	7
17	Evening	2, 3, 4, 5, 6, 7, 8	5
18	Evening	3, 4, 5, 6, 7, 8, 9	5
19	Night	4, 5, 6, 7, 8, 9, 10	6
20	Night	5, 6, 7, 8, 9, 10, 11	5

The overall solution for each unit in the hospital is shown in Table 11. The cell entries indicate what types of nurses were selected to meet the incremental demand requirements. Each resource is denoted by a letter (P = pool, F = float, A = agency, OT = overtime, OC = on-call) and a number. The letters indicate the type of nurse, while the numbers indicate the number of the specific resource needed in the case of agency and on-call nurses, or index of the nurse assigned to the unit in the case of pool and regular nurses.

Table 9. Assignments for regular nurses in problem 1

Regular nurse	Assigned shift	Home unit	Allowed unit	Float to unit	OT unit	Overtime period
1	Off	1	0, 1, 3, 4, 5	-	-	None
2	Day	2	1, 3, 6, 7	6	2	Evening
3	Day	4	3, 5, 6, 7	7	-	None
4	Evening	6	2, 10, 12	2	6	Day
5	Evening	11	3, 11, 12	3	11	Night
6	Night	2	10, 11, 12	12	-	None
7	Night	2	7, 8, 9	8	-	None
8	Off	1	0, 6, 7	-	-	None
9	Day	2	4, 8, 9	8	2	Evening
10	Day	6	5.9, 10	10	-	None
11	Evening	10	8, 9, 10	9	10	Day
12	Evening	10	8, 9, 12	12	10	Late-Day
13	Night	9	8, 9	8	-	None
14	Night	9	9, 10	10	-	None
15	Day	9	0, 4, 5	-	-	None
16	Day	9	0, 4, 7	0	-	None
17	Evening	6	7, 8, 9	8	6	Day
18	Evening	13	6, 7, 8, 9	9	-	None
19	Night	9	10, 11	10	-	None
20	Night	2	10, 11, 13	13	-	None

Table 10. Assignments for agency and on-call nurses for problem 1

	Agency shift			On-call period					
Unit	Day	Evening	Night	1	2	3	4	5	6
0	1	1	1	0	0	0	0	1	1
1	1	1	1	0	0	0	0	1	1
2	0	0	0	0	0	0	0	0	0
3	0	0	0	0	0	0	0	0	0
4	0	0	0	0	0	0	0	0	0
5	1	1	1	0	0	0	0	1	1
6	0	0	0	0	0	0	0	1	1
7	0	0	0	0	0	0	0	1	1
8	1	1	1	0	0	0	0	0	0
9	0	1	0	0	0	0	0	0	0
10	0	0	0	0	0	0	0	1	1
11	0	0	0	0	0	0	0	1	1
12	0	0	0	0	0	0	0	1	1
13	0	0	0	0	0	0	0	1	1

Table 11. Full set of assignments over 24 hours for problem 1

	Period					
Unit	1	2	3	4	5	6
0	F16, 1A	F16, 1A	1A	1A	1OC, 1A	1OC, 1A
1	1A	1A	P11, 1A	P11, 1A	P13, 1OC, 1A	P13, 1OC, 1A
2			F4, OT2, OT9	F4, OT2, OT9		
3	P1	P1	F5, P1	F5, P2	P2	P2
4						
5	1A	1A	P17, P18, 1A	P17, P18, 1A	P14, P20, 1OC, 1A	P14, P20, 1OC, 1A
6	F2, P15, OT4, OT17	F2, P15, OT4, OT17			P6, P7, P19, 1OC	P6, P7, P19, 1OC
7	F3, P16	F3, P16	P4, P5	P4, P5	1OC	1OC
8	F9, P3, 1A	F9, P3, 1A	F17, 1A	F17, 1A	F7, F13, 1A	F7, F13, 1A
9			F11, F18, 1A	F11, F18, 1A		
10	F10, OT11	F10, OT11, OT12			F14, F19, 1OC	F14, F19, 1OC
11	P10	P10			OT5, 1OC	OT5, 1OC
12	P9	P9	F12, P12	F12, P12	F6, 1OC	F6, 1OC
13	P8	P8			F20, 1OC	F20, 1OC

To understand the tradeoff between monetary cost and preference violations, the model was run with different values of P^{max}, the maximum allowable cumulative penalty associated with undesirable patterns. As P^{max} is decreased, costs increase up to the point where the feasible region is empty. Figure 1 depicts the tradeoff curve for problem 1. The point on the far right of the curve corresponds to the unconstrained case whose solution is given in the bottom two rows of Table 7. Here, the cost is \$17,172 and $P^{max} = 258$. As P^{max} is decreased and the model rerun, the objective function increases at an increasing rate until $P^{max} = 80$. At this point, the cost is \$44,022, almost 100% above the best achievable cost. For values of P^{max} below 80, the problem is infeasible. Similar behavior was observed for the other problems investigated.

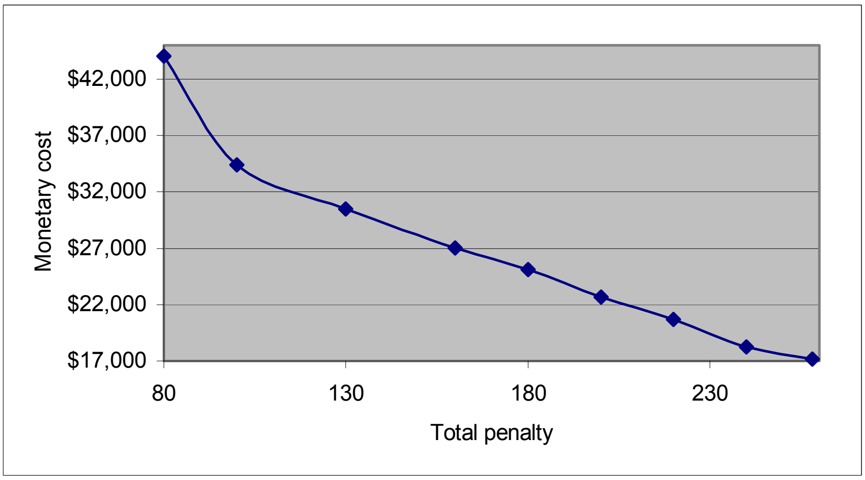

Fig. 1. Cost preference tradeoff for problem 1

6. Discussion

As the focus of management shifts from generating rosters over a 4-week planning horizon to real-time control, the scheduling effort moves from trying to satisfy individual preferences to minimizing costs. The model proposed in this paper allows management to emphasis costs without abandoning the issues that drive midterm scheduling. Because the computational effort is minimal, the methodology allows the user to construct tradeoff curves in a tightly constrained environment.

One of the weaknesses of the model is that it does not allow shifts to be split among units, say, in 4-hour blocks. To remedy this situation, we are now developing a more robust model that offers this feature in a decision support framework. The difficulty is that the size of the decision space grows exponentially with the number of periods over which the regular and pool nurse variables, x and y, are defined. Initially testing of the model indicates that solution times would be on the order of 10 to 15 minutes for problem instances comparable in size to the ones given in Table 1. It is most likely that a decomposition approach would be needed to achieve reasonable efficiency.

References

1. Aickelin, U., Dowsland, K.: Exploiting Problem Structure in a Genetic Algorithm Approach to a Nurse Rostering Problem. J. Sched. 3 (2000) 139-153

2. Bard, J.F., Purnomo, H.W.: Preference Scheduling for Nurses Using Column Generation, Working paper, Graduate Program in Operations Research & Industrial Engineering, The University of Texas, Austin, Texas (2003)
3. Berke, E., Cowling, P., De Causmaecker, P., Vanden Berghe, G.: A Memetic Approach to the Nurse Rostering Problem. Appl. Intel. 15(3) (2001) 199-214
4. Berman, O., Larson, R.C.: A Queueing Control Model for Retail Services Having Back Room Operations and Cross-trained Workers. Comp. Oper. Res. 31 (2004) 201-222
5. Berrada, I., Ferland, J.A., Michelon, P.: A Multi-objective Approach to Nurse Scheduling with both Hard and Soft Constraints. Soc.-Eco. Plan. Sci. 30 (1996) 183-193
6. Cheang, B., Li, H. Lim, A., Rodrigues, B.: Nurse Rostering Problems – A Bibliographic Survey. Euro. J. Oper. Res. 151 (2003) 447-460
7. Jaumard, B., Semet, F., Vovor, T.: A Generalized Linear Programming Model for Nurse Scheduling. Euro. J. Oper. Res. 107 (1998) 1-18
8. Petrovic, S., Beddoe, G., Vanden Berghe, G.: Storing and Adapting Repair Experiences in Employee Rostering. In: Burke, E., De Causmacker, P. (eds.): Practice and Theory of Automated Timetabling IV. Lecture Notes in Computer Science. Vol. 2740. Springer-Verlag, Berlin Heidelberg New York (2003) 148-165
9. Siferd, S.P., Benton, W.: A Decision Modes for Shift Scheduling of Nurses. Euro. J. Oper. Res. 74 (1994) 519-527