Optimal Scheduling of Shifts and Breaks Using Employees Having Limited Time-Availability

Gary M. Thompson

Cornell University

September 1995

Introduction

Labor scheduling is an important function performed by managers of service operations. Customers demand rapid response times, requiring that appropriate numbers of employees are present at all times. Having too many staff on hand lowers operating margins, while having too few servers risks the loss of present or future business. The task of balancing labour capacity and demand – labour scheduling – is made more difficult by demand that often varies widely during the operating day, since the customer-contact activities performed by service delivery personnel cannot be inventoried. Further increasing the difficulty of developing labour schedules is the fact that managers often resort to using employees available for work only during individually specified portions of the operating day (a condition we call limited time-availability). This occurs when labour supplies are tight or when managers desire (or require) the scheduling flexibility offered by employees willing to work part-time.

1

The vast majority of integer and linear programming (LP) labour scheduling models have explicitly represented decision variables (Bailey and Field, 1985; Dantzig, 1954; Henderson and Berry, 1976; Keith, 1979; Krajewski et al., 1980; Mabert and Watts, 1982; Morris and Showalter, 1983; Showalter and Mabert, 1988). In an explicit labour scheduling model, each column contains sufficient information to specify a work schedule exactly. For certain problems, "implicit models" offer advantages. In an implicit model, each column does not contain sufficient information to specify a work schedule exactly. Implicit models thus require post-solution processing of variables to construct the work schedule. Moondra (1976) was the first to define shifts implicitly. Variables represented the number of shifts starting at the beginning and finishing at the end of each period. Constraints imposed minimum and maximum shift durations. Thompson (1990) implicitly matched employees to explicitly-represented shifts. Shifts that could be staffed by the same subset of employees were grouped into sets defined as regions. Employee time-availability was imposed by mapping the employees onto the regions, which were then mapped onto shifts. Bechtold and Jacobs (1990), using constraints to impose bounds on work stretches, matched breaks to shifts implicitly. They found their resultant model to be superior to the standard scheduling formulation of Dantzig (1954). Thompson (1995) integrated Moondra's (1976) implicit representation of shifts with Bechtold and Jacobs' (1990) implicit matching of breaks to shifts, and found the resultant model superior to that of Bechtold and Jacobs (1990). In a recent paper, Jarrah et al. (1994) use Bechtold and Jacobs' (1990) implicit modelling of breaks and an implicit representation of tours to solve large tour scheduling problems. The tour scheduling problem involves identifying both the specific shifts and days on which employees work.

In real service organizations, managers often use part-time employees for the flexibility offered by their short shifts. Part-time employees frequently have non-work (e.g. child care, education) or other work (e.g. second or third job) activities which restrict the times they are available. Such staff can comprise a substantial portion of the workforce in fast-food establishments and supermarkets, for example. Several mathematical-programming-based procedures for scheduling employees having limited time-availability have been developed (Loucks and Jacobs, 1991; Love and Hoey, 1990; Thompson, 1990; Vakharia et al., 1992). Of these, only the models presented by Thompson (1990) both schedule breaks and ensure that all scheduled shifts can be uniquely assigned to individual employees. Shift scheduling models that fail to schedule breaks require that managers assign breaks in real-time. Leaving the assignment of breaks to real-time runs the risk of delivering poorer than desired service (when they are assigned), while employee attitude and performance may decline if breaks are deferred or not assigned [1]. If one ignores the specific time-availability of employees when developing the schedule, one runs the risk of developing a schedule that is unacceptable – that is, not all shifts in the schedule can be uniquely assigned to individual employees. In this event, one would then have to modify the schedule, using some heuristic procedure, or, less desirably, force some employees to work shifts that are in conflict with their other activities.

Shift scheduling flexibility is attained by having many alternative times at which shifts may start or finish, many periods in a shift when a break may be taken, and numerous lengths of shifts. Flexibility, although typically enabling a better match of capacity to demand, also combinatorially complicates the development of schedules. This may explain why the labour scheduling literature largely has allowed only limited flexibility.

Although short duration planning periods have been considered by the labour scheduling literature since the late 1970s (Henderson and Berry, 1976; Keith, 1979), recent increases in computing power have facilitated their use. In general, shorter duration planning periods are preferable under more variable customer demand. Planning periods of ten or 15 minutes also facilitate the scheduling of rest breaks, or reliefs, instead of taking these breaks by managerial or employee discretion in real-time.

In this paper, we present an integer programming (IP) model for scheduling shifts and breaks, which, by directly accounting for the limited time-availability of employees, ensures that all shifts can be uniquely assigned to individual employees. The model is triply-implicit, in that both employees and breaks are implicitly matched to implicitly-defined shifts. We evaluate the new model compared to four other models, including the best of the models presented by Thompson (1990). This evaluation occurs for test environments containing a good deal of shift scheduling flexibility and where planning intervals range from 60 to 15 minutes.

We focus on the shift scheduling problem for several reasons. First, no model has been developed for optimally solving large tour scheduling problems using employees having limited time-availability. Second, no model has been developed for optimally solving high flexibility shift scheduling problems using employees having limited time-availability. Third, optimal solutions to daily shift scheduling problems can be helpful to managers making staffing decisions. For example, by comparing the optimal solution to the shift scheduling problem using employees having limited time-availability to the optimal solution to the shift scheduling problem assuming continuously available employees, one can identify the effect of the limited time-availability. This information, which typically cannot be obtained in a tour model (due to

4

the confounding effect of the across-day restrictions), can be helpful in identifying desired time-availability characteristics of potential employees.

The remainder of the paper describes an environment where employees have limited time-availability; identifies modelling assumptions; illustrates the triply-implicit model by way of an example; presents comparison scheduling procedures and describes the second test environment; and provides results. Finally, we provide a discussion, including suggestions for future research.

An example service environment

In this section we describe our first set of test problems, PS1, which is based on the supermarket cashier scheduling problem described by Thompson (1988). The facility in question operated 15 hours a day, and used planning intervals of 15 minutes. Sixty-three problems, varying on three characteristics, comprise PS1. The first problem characteristic is the pattern of the ideal staffing levels (seven daily patterns). The second problem characteristic is the mean ideal staff size per period. This characteristic has levels of five, ten, and 15 employees. The third characteristic, with three levels, is the degree of scheduling flexibility. Level one, (relatively) low flexibility, allowed shifts of seven to nine hours in duration, including an hour-long meal break. The meal break was preceded and followed by at least three hours of work. Level two, (relatively) intermediate flexibility, allowed all the shifts in level one, but also allowed shifts of 5.5 to 6.5 hours. These shifts contained a half-hour break that was preceded and followed by at least two hours of work. Finally, level three, (relatively) high flexibility, allowed all the shifts of level two, but also shifts of 3 to 4.75 hours. These shifts received no breaks. At all levels, shift length and break timing could vary in increments of 15 minutes. Thus, at level one, shifts of 7, 7.25, 7.50, …, 8.50, 8.75, and 9.00 hours would be acceptable. For a nine-hour shift, the meal

5

break could be taken after 3, 3.25, …, 4.75, or 5 hours of work. Given the restrictions defining acceptable shifts, there are 1,245, 2,530, and 2,894 different shifts (defined by start time, number of working periods, and break timing) that could be scheduled with the three levels of scheduling flexibility, ignoring the time-availability of employees.

Figure 1 illustrates the ideal staff sizes, by period, for the seven days with a mean ideal staff size of 15 employees per period. The employee time-availability data for each day of the week were actual data from an organization (Thompson, 1988, Table II, p. 238). The days had from 30 to 34 available employees, and averaged 31.6 employees available.

Modelling assumptions

We make the following assumptions in this study:

(1) Each employee is available for work in an individually specified subset of the operating day.

(2) For each planning period, the desired staff size has been predetermined using an appropriate methodology.

(3) Other than specified in the definition of shifts, there is no minimum acceptable shift duration for any employee.

(4) There are no split availabilities (that is, there are no employees who are available for work during two or more non-contiguous time stretches)[2].

(5) All employees have identical skills.

(6) Break durations are fixed.

(7) Shifts receive at most one scheduled break.

We note that assumption one is key to the current investigation. Assumption one is not a limiting assumption, but rather serves to extend the range of conditions typically considered in the scheduling literature. Indeed, the common assumption of continuously available employees is simply a special case of assumption one.

The new, triply-implicit model (ISIBIM)

We refer to the new model as ISIBIM, since it implicitly represents shifts, it implicitly represents breaks, and it implicitly matches employees to shifts. Our motivations for developing ISIBIM were the difficulty of optimal shift scheduling using employees having limited time-availability and the fact that implicit models have performed exceedingly well compared to explicit models (Bechtold and Jacobs, 1990; Jarrah et al., 1994; Thompson, 1990, 1995). Our goal became the alteration of Thompson's (1995) model for the environment of employees having limited time-availability. Several characteristics of that model made the extension non-trivial. First, it implicitly represents shifts. That is, the only information about shifts in his model is the number of shifts starting in each period, the number of shifts finishing in each period, and the number of shifts commencing a break in each period. Characteristics like the acceptable length of shifts and the timing of breaks are imposed via constraint sets. With the implicit representation of shifts, mapping the time-availability of individual employees onto the shifts proved difficult. For a mapping to be effective, any feasible solution to the new model should be availability-feasible, that is, one must be able to assign all scheduled shifts uniquely to individual employees. In particular, the condition that caused the greatest problem was one we call extraordinary time-availability overlap, or ETAO, which occurs when one employee is available both earlier and later than another employee. The key insight necessary for extending Thompson's (1995) model to the environment of limited time-availability was separating the

7

employees into groups, such that ETAO did not exist within any group. Then, the mapping of employees onto shifts was handled separately for each availability group.

We illustrate ISIBIM using a simple scheduling problem. Assume that the operating day is 12 hours and that acceptable shifts are of seven to nine hours in duration. These shifts include an hour-long meal break which must be preceded and followed by at least two hours of work. Further, assume that six employees constitute the workforce and that these employees are available during the operating day as shown in Table I. Finally, assume that 1, 3, 5, 5, 5, 5, 5, 5, 5, 5, 3, and 1 employees are desired in periods 1-12, respectively.

If all employees were continuously available for work, the optimal solution to this problem would be to schedule the six shifts identified in Table II. This schedule provides the ideal number of staff in each period. Unfortunately, however, when one considers the actual availability of the employees (as identified in Table I) the schedule is not availability-feasible. Specifically, shifts 5 and 6 would require employees 5 and 6 to work beyond the times they are available.

ISIBIM ensures that any feasible solutions to it are availability-feasible by representing the time-availability of employees within the model. As noted earlier, we first divide the employees into groups, such that no ETAO occurs within any group. Given the availability data presented in Table I, two occurrences of ETAO exist: employee 1 is available both earlier and later than employee 5, and employee 3 is available both earlier and later than employee 6. Thus, we separate the employees into two availability-based groups, with group one comprising employees 1-4, and group two comprising employees 5 and 6. No ETAO then exists within either group.

8

The first part of ISIBIM is similar to Thompson's model (1995), except that ISIBIM defines separate variable sets for each availability-based group. We thus have the variables $s_{1,6}$, ..., $s_{1,6}$, representing the number of availability group-one shifts that begin in periods 1-6, respectively and $s_{2,2}$, ..., $s_{2,5}$, representing the number of availability group-two shifts that begin in periods 2-5 respectively (note that no one in group 2 can work shifts that begin in periods 1 or 7). The variables $f_{1,7}$, ..., $f_{1,12}$, represent the number of availability group-one shifts that finish at the end of periods 7-12 respectively, with $f_{2,8}$, ..., $f_{2,11}$ representing the number of availability group-two shifts that finish at the end of hours 8-11 respectively (again, no one in group 2 can work shifts that finish in hours 7 or 12). Define the variables b_3, ..., b_{10} as the number of employees who commence a meal break in hours 3-10 respectively.

ISIBIM's objective is to minimize the net cost of the schedule:

$$Min\ Z = 7f_{1,7} + 8f_{1,8} + 9f_{1,9} + 10f_{1,10} + 11f_{1,11} + 12f_{1,12} + 9f_{2,9} + 10f_{2,10} + 11f_{2,11} - 1s_{1,2}$$
$$- 2s_{1,3} - 3s_{1,4} - 4s_{1,5} - 5s_{1,6} - 1s_{2,2} - 2s_{2,3} - 3s_{2,4} - 4s_{2.6} - b_3 - b_4 - b_5$$
$$- b_6 - b_7 - b_8 - b_9 - b_{10} + M \sum_{p=1}^{12} \mu.$$

The objective has two components. The first component, $7f_{1,7}, 7, ..., -b_{10}$, represents the total scheduled employee-periods. The second component, $M \sum_{p=1}^{12} \mu_p$, represents the artificial penalty associated with being understaffed, i.e. having fewer than the desired number of staff present in any period. Although understaffing is undesirable, it may be unavoidable due to the limited time-availability of employees.

ISIBIM's first set of restrictions measures the staffing levels in each period compared to the desired staff size. The number of staff working in a period is given by the number of shifts

that have commenced in the period or earlier, less the number of shifts that have finished work before the current period, and less the number of employees on break during the period. For each period, ISIBIM also uses a general integer variable to measure any shortfall of staff compared to the desire staff size. Thus:

$$s_{1,1} + \mu_1 \geq 1 \text{ (period 1)}$$

$$s_{1,1} + s_{1,2} + s_{2,2} + \mu_2 \geq 3 \text{ (period 2)}$$

$$s_{1,1} + s_{1,2} + s_{1,3} + s_{2,2} + s_{2,3} - b_3 + \mu_3 \geq 5 \text{ (period 3)}$$

$$s_{1,1} + \cdots + s_{1,4} + s_{2,2} + s_{2,3} + s_{2,4} - b_4 + \mu_4 \geq 5 \text{ (period 4)}$$

$$s_{1,1} + \cdots + s_{1,5} + s_{2,2} + \cdots + s_{2,5} - b_5 + \mu_5 \geq 5 \text{ (period 5)}$$

$$s_{1,1} + \cdots s_{1,6} + s_{2,2} + \cdots + s_{2,5} - b_6 + \mu_6 \geq 5 \text{ (period 6)}$$

$$s_{1,1} + \cdots + s_{1,6} + s_{2,2} + \cdots + s_{2,5} - b_7 + \mu_7 \geq 5 \text{ (period 7)}$$

$$s_{1,1} + \cdots + s_{1,6} + s_{2,2} + \cdots + s_{2,5} - f_{1,7} - b_8 + \mu_8 \geq 5 \text{ (period 8)}$$

$$s_{1,1} + \cdots + s_{1,6} + s_{2,2} + \cdots + s_{2,5} - f_{1,7} - f_{1,8} - f_{2,8} - b_9 + \mu_9 \geq 5 \text{ (period 9)}$$

$$s_{1,1} + \cdots + s_{1,6} + s_{2,2} + s_{2,5} - f_{1,7} - f_{1,8} - f_{2,8} - f_{2,9} - b_{10} + \mu_{10} \geq 5 \text{ (period 10)}$$

$$s_{1,1} + \cdots + s_{1,6} + s_{2,2} + \cdots + s_{2,5} - f_{1,7} - \cdots - f_{1,10} - f_{2,8} - f_{2,9} - f_{2,10} + \mu_{11} \geq$$
$$3 \text{ (period 11)}$$

$$s_{1,1} + \cdots + s_{1,6} + s_{2,2} + \cdots + s_{2,5} - f_{1,7} - \cdots - f_{1,11} - f_{2,8} - \cdots - f_{2,11} + \mu_{12} \geq$$
$$1 \text{ (period 12)}$$

The next set of constraints equates the number of shift starts to the number of shift finishes, by availability group:

$$s_{1,1} + \cdots + s_{1,6} = f_{1,7} + \cdots + f_{1,12}$$

$$s_{2,2} + \cdots + s_{2,5} = f_{2,8} + \cdots + f_{2,11}$$

Because of its implicit representation of shifts, ISIBIM includes constraints that impose the minimum and maximum acceptable shift durations. First, the minimum acceptable durations are imposed for each availability group:

$$s_{1,1} \geq f_{1,7}$$

$$s_{1,1} + s_{1,2} \geq s_{1,3} \geq f_{1,7} + f_{1,8} + f_{1,9}$$

$$s_{1,1} + s_{1,2} + s_{1,3} \geq f_{1,7} + f_{1,8} + f_{1,9}$$

$$s_{1,1} + \cdots + s_{1,4} \geq f_{1,7} + \cdots + f_{1,10}$$

$$s_{1,1} + \cdots + s_{1,5} \geq f_{1,7} + \cdots + f_{1,11}$$

$$s_{2,2} \geq f_{2,8}$$

$$s_{2,2} + s_{2,3} \geq f_{2,8} + f_{2,9}$$

$$s_{2,2} + s_{2,3} + s_{2,4} \geq f_{2,8} + f_{2,9} + f_{2,10}$$

Then, the maximum acceptable durations are imposed for each availability group. Note that such restrictions are not necessary for availability group two, since the length of shifts is limited by the restrictions (below) which impose the individual time-availabilities. Thus:

$$s_{1,1} \leq f_{1,7} + f_{1,8} + f_{1,9}$$

$$s_{1,1} + s_{1,2} \leq f_{1,7} + \cdots + f_{1,10}$$

$$s_{1,1} + s_{1,2} + s_{1,3} \leq f_{1,7} + \cdots + f_{1,11}$$

Another restriction equates the number of shift starts to the number of breaks scheduled:

$$s_{1,1} + \cdots + s_{1,6} + s_{2,2} + \cdots s_{2,5} = b_3 + \cdots + b_{10}$$

Then, a set of restrictions enforces the minimum of two hours of work prior to a break:

$$s_{1,1} \geq b_3$$

$$s_{1,1} + s_{1,2} + s_{2,2} \geq b_3 + b_4$$

$$s_{1,1} + s_{1,2} + s_{1,3} + s_{2,2} + s_{2,3} \geq b_3 + b_4 + b_5$$

$$s_{1,1} + \cdots + s_{1,4} + s_{2,2} + s_{2,3} + s_{2,4} \geq b_3 + \cdots + b_6$$

$$s_{1,1} + \cdots + s_{1,5} + s_{2,2} + \cdots + s_{2,5} \geq b_3 + \cdots + b_7$$

Another set of restrictions enforces the minimum of two hours of work after a break:

$$f_{1,12} \geq b_{10}$$

$$f_{1,11} + f_{1,12} + f_{2,11} \geq b_9 + b_{10}$$

$$f_{1,10} + f_{1,11} + f_{1,12} + f_{2,10} + f_{2,11} \geq b_8 + b_9 + b_{10}$$

$$f_{1,9} + \cdots + f_{1,12} + f_{2,9} + f_{2,10} + f_{2,11} \geq b_7 + \cdots + b_{10}$$

$$f_{1,8} + \cdots + f_{1,12} + f_{2,8} + \cdots + f_{2,11} \geq b_6 + \cdots + b_{10}$$

To this point, ISIBIM has largely resembled the model of Thompson (1995), except for its use of variables and constraints for each availability-based group. The remainder of ISIBIM includes

new variables and constraints necessary for imposing the limited time-availability of individual employees. The new variables serve to indicate whether individual employees work, while the new constraints limit the number and times at which shifts can start and finish based on the employees' time-availability. Define $k_1, \ldots, k_6$ as taking a value of 1 if employees 1-6 respectively are scheduled to work, and a value of zero otherwise. The first new constraint set equates the number of shifts scheduled from each availability group to the number of employees from that group who actually work. Thus:

$$s_{1,1} + \cdots + s_{1,6} = k_1 + k_2 + k_3 + k_4$$

$$s_{2,2} + \cdots + s_{2,5} = k_5 + k_6$$

ISIBIM then restricts the number and times at which shifts from each availability group may start. First consider availability group 1. Employees 1 and 2 in group 1 become available for work in the first hour, while employees 3 and 4 do not become available until hour 3. Thus, if any shifts start prior to hour three (i.e., in hour 1 or hour 2) they must be staffed by either or both of employees 1 or 2:

$$s_{1,1} + s_{1,2} \leq k_1 + k_2$$

For availability group 2, any shift that starts before hour 4 (i.e., in hour 2 or hour 3) must be staffed by employee 5:

$$s_{2,2} + s_{2,3} \leq k_5$$

For each availability group, ISIBIM then restricts the number and times at which shifts can finish. For availability group 1, employee 3 can finish as late as hour 12, employees 1 and 4 can finish as late as hour 11, and employee 2 can finish no later than hour 9. Thus, if a shift finishes

in period 12, it must be staffed by employee 3, while those shifts that finish later than hour 9 (i.e. in hours 10, 11, or 12) must be staffed by employees 1, 3, or 4:

$$f_{1,12} \leq k_3$$

$$f_{1.10} + f_{1,11} + f_{1,12} \leq k_1 + k_2 + k_3$$

For availability group 2, any shift ending after hour 9 must be staffed by employee 6:

$$f_{2,10} + f_{2,11} \leq k_6$$

Table III reports the optimal solution to ISIBIM for this problem. Matching shift starts to finishes, and break starts to shifts gives the optimal work schedule shown in Table IV. Note that this schedule only schedules four employees in each of periods 3 and 10, compared to the desired number of five in these periods. As such, the optimal schedule, considering the employees' limited time-availability, is 2M–2 labour hours more costly than the optimal schedule that assumes all employees are continuously available. The higher cost arises from the shortage penalty for two employee periods (2M), less the two fewer labour hours (–2).

Having shown the variables and constraints used to impose the limited availability of employees, it is now possible to discuss the necessity of splitting employees into groups such that no ETAO exists within groups. Consider only employees 1 and 5, and two shifts, one of which covers periods 1-7, the other covering periods 4-10 (ignore breaks). Both shifts, on their own, are of acceptable duration, and each, on its own, can be assigned to an employee. Now assume that employees 1 and 5 were not separated into two groups (i.e. that there are no "availability groups"). Since only one shift starts earlier than hour 2, and only one shift finishes later than hour 9, and since there are two shifts and two employees, all the restrictions provided above would be satisfied. Despite this, the shifts are not both staffable. Problems like this can

14

only be avoided by separating the employees into groups such that no ETAO exists within any group, as ISIBIM does.

Comparison procedures and test environments

In this section, we present the comparison scheduling procedures. We also describe the second problem set, PS2, used to evaluate the effectiveness of ISIBIM relative to other models for scheduling employees having limited time-availability. We developed PS2 to present a wider range of characteristics than that presented in PS1.

Comparison models

In total, we evaluate five models, including ISIBIM. As shown in Table V, which summarizes them, three of the models, including ISIBIM, have never before been evaluated. The models vary in their degree of implicitness. ESEBEM is the most explicit model, while ISIBIM is the most implicit model.

For the test problems that we describe below, each model was generated using GAMS (Brooke et al., 1992) and solved using OSL (IBM Corporation, 1991), and its solution was interpreted and subsequent steps were performed by a FORTRAN program. We report schedule generation times on a Pentium 90-based personal computer. This computer had 32 megabytes of RAM (random access memory), almost all of which was available to GAMS. When solving the models, we limited OSL to using no more than 28 megabytes of RAM.

15

Problem set 2 (PS2)

In developing PS2 our goal was to create a set of test problems more broadly representative of those existing in service organizations who employ people with limited time-availability. To do this, we developed a set of 432 test problems, varying on six dimensions:

(1) the planning interval duration (three levels);

(2) the pattern of ideal staffing levels (four levels);

(3) the mean ideal staffing levels (three levels);

(4) the length of the operating day (three levels);

(5) the degree of scheduling flexibility (two levels); and

(6) the degree of employee time-availability overlap (two levels).

PS2 used planning intervals of 60, 30, and 15 minutes. In developing the ideal staffing level patterns, we ensured that the test problems with shorter duration periods also exhibit greater variability in the ideal staffing levels. One justification for shorter duration periods is greater variability in the customer arrival rate (and thus staffing levels) across the operating day. The test problems with 15-, 30-, and 60-minute planning intervals have coefficients of variation in the ideal staffing levels of approximately 0.6, 0.4, and 0.2 respectively.

We used four types of ideal staffing level patterns: unimodal, bimodal, trimodal and random. Unimodal patterns, with a single mid-day peak, are commonly observed in retail facilities on weekends. Bimodal patterns, with two peaks in daily requirements, often occur in organizations where commuters constitute a majority of customers. Trimodal patterns, with three peaks in daily requirements, commonly occur in restaurants open for the three main meal periods. Random (but predictable) patterns, with several or many peaks in daily demand, also

16

occur in organizations (for example, in a telemarketing firm which handles calls from several time zones).

All of the ideal staffing level patterns were developed with mean ideal staffing levels of five, ten and 15 employees per period. The low end of this staffing range may be observed in medium-sized banks, while the higher end can be seen in large supermarkets. Three lengths of operating day were selected: 14, 18 and 22 hours. The combinations of the various durations of planning intervals, ideal staffing level patterns, mean ideal staffing levels and operating day lengths resulted in a total of 108 ideal staffing level curves. Figure 2 illustrates the 24 curves for the 18-hour operating day and 60- and 15-minute planning intervals.

To investigate the flexibility given by increasing the number of shift alternatives, PS2 used two levels of scheduling flexibility. For both levels, shifts of six to eight working hours, each assigned an hour-long break, were allowed. For the (relatively) high level of scheduling flexibility, split shifts – shifts of six to eight working hours, each assigned a four-hour break – could also be scheduled. The restrictions on break timing were identical for all shifts: a break had to be preceded and followed by at least two, but no more than six, hours of work. A substantial degree of scheduling flexibility is represented by the range of alternate shift durations, the timing of shift starts, and break timing. Table VI shows that, ignoring the employees' limited time-availability, the problems range from 82 shift alternatives, with the low level of scheduling flexibility and 60-minute planning intervals, to 11,814 shift alternatives, with the high level of scheduling flexibility and 15-minute planning intervals.

Varying the minimum required number of availability groups was desirable because the number of variables and constraints in ISIBIM is influenced by the number of such groups. The number of availability groups required – one and two for the low and high levels of this factor

17

respectively [3] – was controlled during the generation of the test data by assigning different availability durations for individual employees. With the single availability group, all employees were available for ten hours for the 14-hour day, 13 hours for the 18-hour day, and 16 hours for the 22-hour day. In contrast, with two availability groups, the employees were separated into approximately equal groups, having availability of eight and 12 hours for the 14-hour day, ten and 16 hours for the 18-hour day, and 12 and 20 hours for the 22-hour day.

The number of individual employees in the PS2 problems ranged from a low of 12, with 14-hour days and a mean ideal staffing level of five employees per period, to a high of 58 employees, with the 22-hour day and a mean ideal staffing level of 15 employees per period. On average, the test problems had 31.3 individual employees available for work. In generating the availabilities we attempted to keep the number of employees to as small a number as might reasonably be able to cover the desired staffing levels.

Results

Table VII reports results on PS1. The table contains data on the size of the model realizations and the time required to generate optimal solutions. We measured the size of the model realizations in two ways: the total number of elements and the number of non-zero elements in their A-matrices. On average, realizations of ESEBEM had the most non-zero A-matrix elements, and the highest schedule generation times. Realizations of ISIBIM had the least number of non-zero A-matrix elements, and the lowest schedule generation times.

Table VIII summarizes the solution times for PS2. In ascending order of schedule generation times the models were ISIBIM, ESIBIM, ESEBIM, ESIBEM, and ESEBEM. Many

of the problems were too large to solve when modelled using ESEBEM, ESIBEM or ESEBIM.

Only ESIBIM and ISIBIM yielded solutions to all PS2 problems.

Discussion

Our discussion focuses on choosing the best model, providing recommendations for future research, and conclusions.

Selecting the best model

The schedule generation time results illustrate the benefit of implicit modelling. The explicit model ESEBEM was generally the largest and took the longest to solve to optimality. Indeed, its size was such that for many problems we were unable even to generate the complete model, let alone solve it, on our computer (despite having 32 megabytes of RAM). The two models with one implicit characteristic, ESEBIM and ESIBEM, each had similar and high schedule generation times (though lower than ESEBEM). The formulation with an intermediate amount of implicit representation, ESIBIM, had an intermediate schedule generation time. The most implicit model, ISIBIM, had the lowest average schedule generation time. ISIBIM is effective in succinctly representing problems with the correct structure. Flexibility, often the bane of scheduling research, increases the value of ISIBIM, since it is this flexibility that the totally implicit approach handles best. For example, shifts of additional durations can be incorporated into ISIBIM with a few additional constraints and only a small fraction of the variables required by other models. The flexibility which ISIBIM models best is short planning intervals, contiguous and numerous shift lengths, large and regular break windows, abundant periods in which shifts may start and/or finish, high average employee availability, and low availability overlap. As Table VI shows, the number of shift alternatives in PS2 is very large,

19

particularly with the shorter planning intervals, the higher level of scheduling flexibility, and longer operating days. These problems would probably be hard to solve optimally, even without the complication of the limited time-availability of employees (Thompson, 1995).

Given these findings, it is rather easy to recommend a model: one should select the model with the greatest degree of implicit representation that allows an adequate portrayal of the true problem. Implicit modelling is effective in lowering schedule generation times because the number of variables is greatly reduced at the expense of a moderate increase in the number of constraints.

Research extensions

Because of the overall superiority of ISIBIM, we address our comments to that model. First, ISIBIM can be extended to continuous (24-hour) operations. In this event, one may have to define additional variables to ensure that the model "wrapped around" correctly. Second, given that employee time-availability frequently varies across days in many service environments, disaggregation approaches to tour scheduling – whereby weekly schedules are developed by combining separate daily schedules generated with ISIBIM – may have more potential than in environments where employees are continuously available. Third, the effectiveness of ISIBIM may serve to stimulate the development of implicit models for other scheduling environments. One promising, though undoubtedly challenging, area is the scheduling of employees having heterogeneous skills. Fourth, developing heuristics based on solutions to the LP relaxation of ISIBIM may also be fruitful. Finally, there is a need for work addressing problems where there are large numbers of employees (several hundred, say) having limited time-availability. It may be possible to ignore the limited time-availability during the development of the schedule, yet

deal with it acceptably during the assignment of individuals to shifts. This approach may be required if ISIBIM's binary variables noticeably slow its solution in such problems.

Conclusions

This paper has presented a new, triply-implicit model for the scheduling of shifts and breaks using employees who are available for work only during individually-specified portions of the operating day. The new model performed better than all existing comparable models. Specifically, it yielded optimal solutions markedly faster than the existing comparable models, particularly with shorter planning intervals and more alternate shifts, both conditions desirable to managers who wish to match actual staffing levels to their ideal level as closely as possible.

Notes

1. We concur with one of the referees that forecast accuracy should play a role in break scheduling. When forecasts of customer demand are reasonably accurate, breaks should probably be scheduled in advance, and taken when scheduled. When forecasts of demand are less accurate, breaks should still be scheduled in advance, but rescheduled in real-time.

2. This assumption can be relaxed in our model by adding a constraint for each employee who is available in multiple, non-continuous time stretches. The constraint would restrict the employee to working in only one of his/her available time stretches.

3. In PS1, one of the days required only one availability-based group, five days required two groups, and one day required three groups.

References

Bailey, J. and Field, J. (1985), "Personnel scheduling with flexshift models", *Journal of Operations Management*, Vol. 5 No. 3, pp. 327-38.

Bechtold, S.E. and Jacobs, L.W. (1990), "Implicit optimal modeling of flexible break assignments in labor staffing decisions for service operations", *Management Science*, Vol. 36 No. 11, pp. 1339-50.

Brooke, A., Kendrick, D. and Meeraus, A. (1992), *GAMS, Release 2.25, A User's Guide*, The Scientific Press, South San Francisco, CA.

Dantzig, G.B. (1954), "A comment on Edie's 'Traffic delays at toll booths'", *Operations Research*, Vol. 2 No. 3, pp. 339-41.

Henderson, W.B. and Berry, W.L. (1976), "Heuristic methods for telephone operator shift scheduling: an experimental analysis", *Management Science*, Vol. 22 No. 12, pp. 1372-80.

IBM Corporation (1991), *Optimization Subroutine Library, Release 2*, IBM Corporation, Kingston, NY.

Jarrah, A.I.Z., Bard, J.F. and deSilva, A.H. (1994), "Solving large-scale tour scheduling problems", *Management Science*, Vol. 40 No. 9, pp. 1124-44.

Keith, E.G. (1979), "Operator scheduling", *AIIE Transactions*, Vol. 11 No. 1, pp. 37-41.

Krajewski, L.J., Ritzman, L.P. and McKensie, P. (1980), "Shift scheduling in banking operations: a case application", *Interfaces*, Vol. 10 No. 2, pp. 1-8.

Loucks, J.S. and Jacobs, F.R. (1991), "Tour scheduling and task assignment of a heterogeneous work force: a heuristic approach", *Decision Sciences*, Vol. 22 No. 4, pp. 719-38.

Love, R.R. Jr and Hoey, J.M. (1990), "Management science improves fast-food operations", *Interfaces*, Vol. 20 No. 2, pp. 21-9.

Mabert, V.A. and Watts, C.A. (1982), "A simulation analysis of tour-shift construction procedures", *Management Science*, Vol. 28 No. 5, pp. 520-32.

Moondra, S.L. (1976), "An L. P. model for work force scheduling for banks", *Journal of Bank Research*, Vol. 7 No. 4, pp. 299-301.

Morris, J.G. and Showalter, M.J. (1983), "Simple approaches to shift, days-off and tour scheduling problems", *Management Science*, Vol. 29 No. 8, pp. 942-50.

Showalter, M.J. and Mabert, V.A. (1988), "An evaluation of a full-/part-time tour scheduling methodology", *International Journal of Operations and Production Management*, Vol. 8 No. 7, pp. 54-71.

Thompson, G.M. (1988), "A comparison of techniques for scheduling non-homogeneous employees in a service environment subject to non-cyclical demand", unpublished PhD dissertation, The Florida State University, Tahallassee, FL.

Thompson, G.M. (1990), "Shift scheduling when employees have limited availability: an L.P. approach", *Journal of Operations Management*, Vol. 9 No. 3, pp. 352-70.

Thompson, G.M. (1995), "Improved implicit optimal modeling of the shift scheduling problem", *Management Science*, Vol. 41 No. 4, pp. 595-607.

Vakharia, A.J., Selim, H.S. and Husted, R.R. (1992), "Efficient scheduling of part-time employees", *Omega*, Vol. 20 No. 2, pp. 201-13.

Table 1. Availability of the employees in the sample problem

Employee	Earliest possible start time (start of hour)	Latest possible finish time (end of hour)
1	1	11
2	1	9
3	3	12
4	3	11
5	2	9
6	4	11

Table 2. An optimal schedule to the sample problem, assuming all employees are continuously available.

Period	Shift 1	2	3	4	5	6	Total working	Ideal staff size	Surplus/ (shortage)
1	W						1	1	–
2	W	W			W		3	3	–
3	W	W	W		W	W	5	5	–
4	B	W	W	W	W	W	5	5	–
5	W	B	W	W	W	W	5	5	–
6	W	W	W	W	B	W	5	5	–
7	W	W	B	W	W	W	5	5	–
8	W	W	W	W	W	B	5	5	–
9	W	W	W	B	W	W	5	5	–
10		W	W	W	W	W	5	5	–
11			W	W		W	3	3	–
12				W			1	1	–

Notes:
W = working; B = break

Table 3. Optimal solution to ISIBIM.

Period (p)	Start, finish and break variable values					Employee working variables	Value
	$s_{1,p}$	$f_{1,p}$	$s_{2,p}$	$f_{2,p}$	b_p		
1	1					k_1	1
2	1		1			k_2	1
3	1					k_3	1
4	1		1		1	k_4	1
5					1	k_5	1
6					1	k_6	1
7					1		
8					1		
9		1		1	1		
10		1					
11		1		1			
12		1					

27

Table 4. An optimal schedule to the sample problem, considering the time-availability of employees.

Period	Shift 1	2	3	4	5	6	Total working	Ideal staff size	Surplus/(shortage)
1	W						1	1	0
2	W	W			W		3	3	0
3	W	W	W		W		4	5	(1)
4	B	W	W	W	W	W	5	5	0
5	W	B	W	W	W	W	5	5	0
6	W	W	W	W	B	W	5	5	0
7	W	W	B	W	W	W	5	5	0
8	W	W	W	W	W	B	5	5	0
9	W	W	W	B	W	W	5	5	0
10		W	W	W		W	4	5	(1)
11		W	W			W	3	3	0
12				W			1	1	0

Notes: Reconstructed from the optimal solution presented in Table III
W = working
B = break

28

Table 5. A summary of the models evaluated, their structure, and origin.

Model	Shift definition	Break definition	Matching employees to shifts	Model origin	Previously evaluated
ESEBEM	Explicit	Explicit	Explicit	Thompson (1990)	Yes
ESIBEM	Explicit	Implicit	Explicit	An extension of Bechtold and Jacobs' (1990) model to explicitly account for employees having limited time-availability	No
ESEBIM	Explicit	Explicit	Implicit	Thompson (1990)	Yes
ESIBIM	Explicit	Implicit	Implicit	An integration of the implicit matching of breaks to shifts, introduced by Bechtold and Jacobs (1990), with the implicit matching of employees to shifts, introduced by Thompson (1990)	No
ISIBIM	Implicit	Implicit	Implicit	Presented here for the first time	No

Table 6. Number of unique shifts in problem set 2, ignoring availability feasibility.

Planning interval duration (minutes)	Level of scheduling flexibility					
	Low (relatively)			High (relatively)		
	Length of operating day (hours)					
	14	18	22	14	18	22
60	82	130	178	128	224	320
30	445	725	1,005	680	1,240	1,800
15	2,865	4,737	6,609	4,326	8,070	11,814

Table 7. Results on problem set 1.

Performance criterion	Number of shift types	ESEBEM	ESIBEM	IP model ESEBIM	ESIBIM	ISIBIM
Number of	1	0.257E + 06	76.94	62.27	59.78	65.72
elements in the	2	0.637E + 06	78.61	45.99	49.56	57.89
A-matrix[a]	3	0.806E + 06	118.64	47.04	56.02	73.48
Number of	1	0.957E + 05	53.51	41.94	21.91	8.21
non-zeros in the	2	0.218E + 06	49.72	32.75	16.91	7.92
A-matrix[b]	3	0.253E + 06	56.63	30.73	17.14	10.87
Solution	1		148.53	451.54	125.95	117.85
time[c]	2		442.54	655.29	156.61	138.08
	3		857.64		211.04	158.58

Notes: All values are averages for 21 problems in each category (seven days multiplied by three mean ideal staffing levels)

[a]A-matrix size for ESEBEM is the number of variables multiplied by the number of constraints. For the other models, size is measured as a percentage of the size of ESEBEM

[c]bValues are the actual number of non-zeros for ESEBEM. For the other models, values are the number of non-zeros as a percentage of the number of non-zeros in ESEBEM

[c]Clock seconds on a Pentium 90-based personal computer with the models generated using GAMS (Brooke *et al.*, 1992), solved using OSL (IBM Corporation, 1991) and results interpreted using FORTRAN programs. Missing values are for problems that were not solved

Table 8. Summary of schedule generation time results for problem set 2[a]

Problem group	ESEBEM	ESIBEM	IP model ESEBIM	ESIBIM	ISIBIM
183 problems solved by all models	342.07	152.49	133.61	114.37	113.29
57 problems solved by all models except ESEBEM	N/A	431.52	453.88	123.32	141.18
70 problems solved only by ESEBIM, ESIBIM and ISIBIM	N/A	N/A	534.94	145.86	182.71
122 problems solved only by ESIBIM and ISIBIM	N/A	N/A	N/A	880.57	416.67
Average for all problems[d]	N/A	N/A	N/A	337.03[b]	213.89[c]

Notes: Average clock seconds on a Pentium 90-based personal computer, with the models generated using GAMS (Brooke *et al.*, 1992), solved using OSL (IBM Corporation, 1991) and results interpreted using FORTRAN programs

[a] The difference between solution times is significant at the $\alpha = 0.01$ level

[b] Includes nine problems not solved to verified-optimality

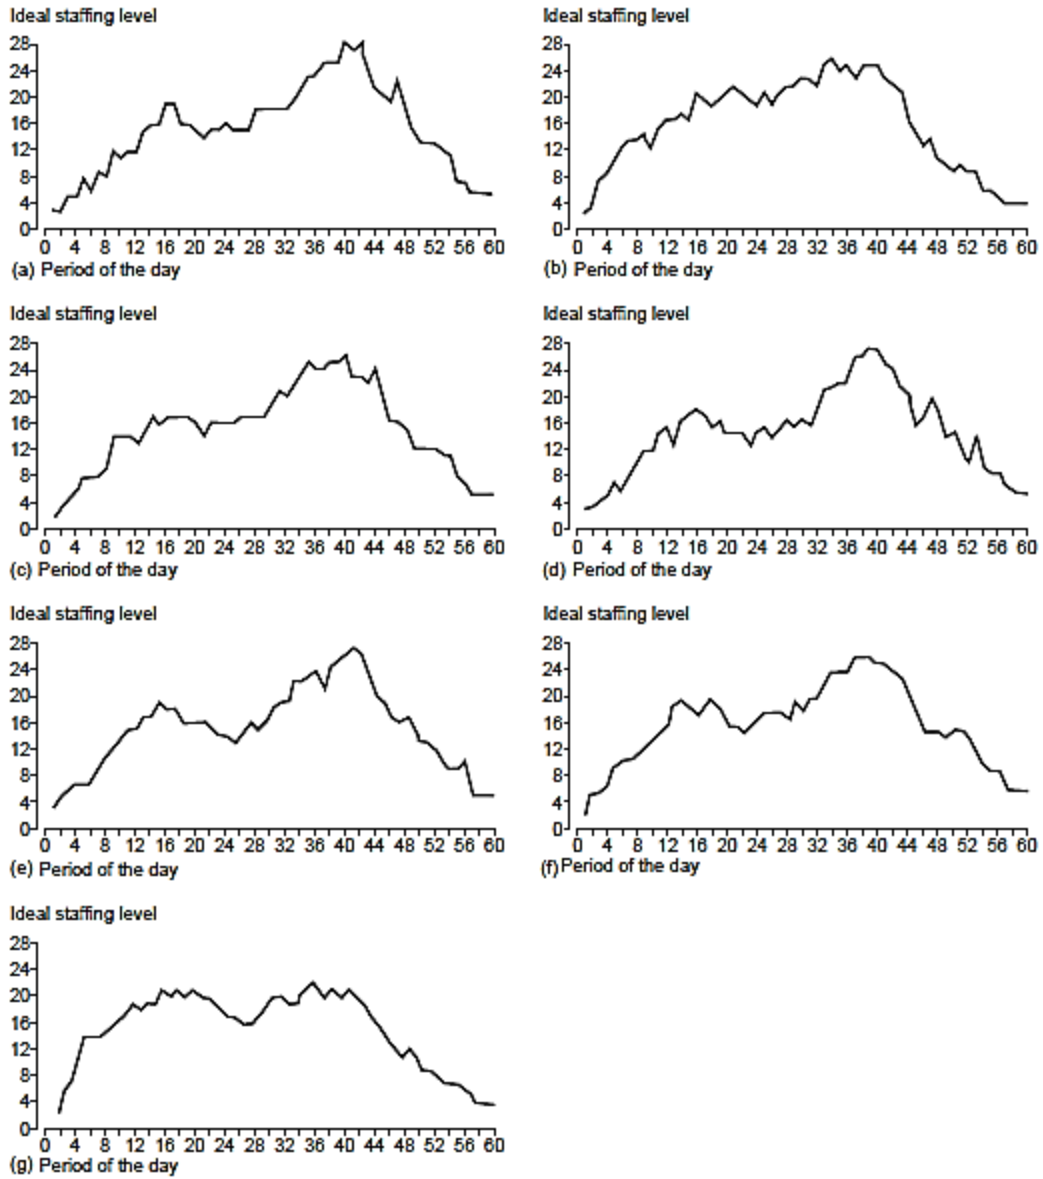[c] Includes five problems not solved to verified-optimality

Figure 1. Examples of ideal staffing levels in problem set 1 (shown for the mean ideal staffing levels of 15 employees per period).
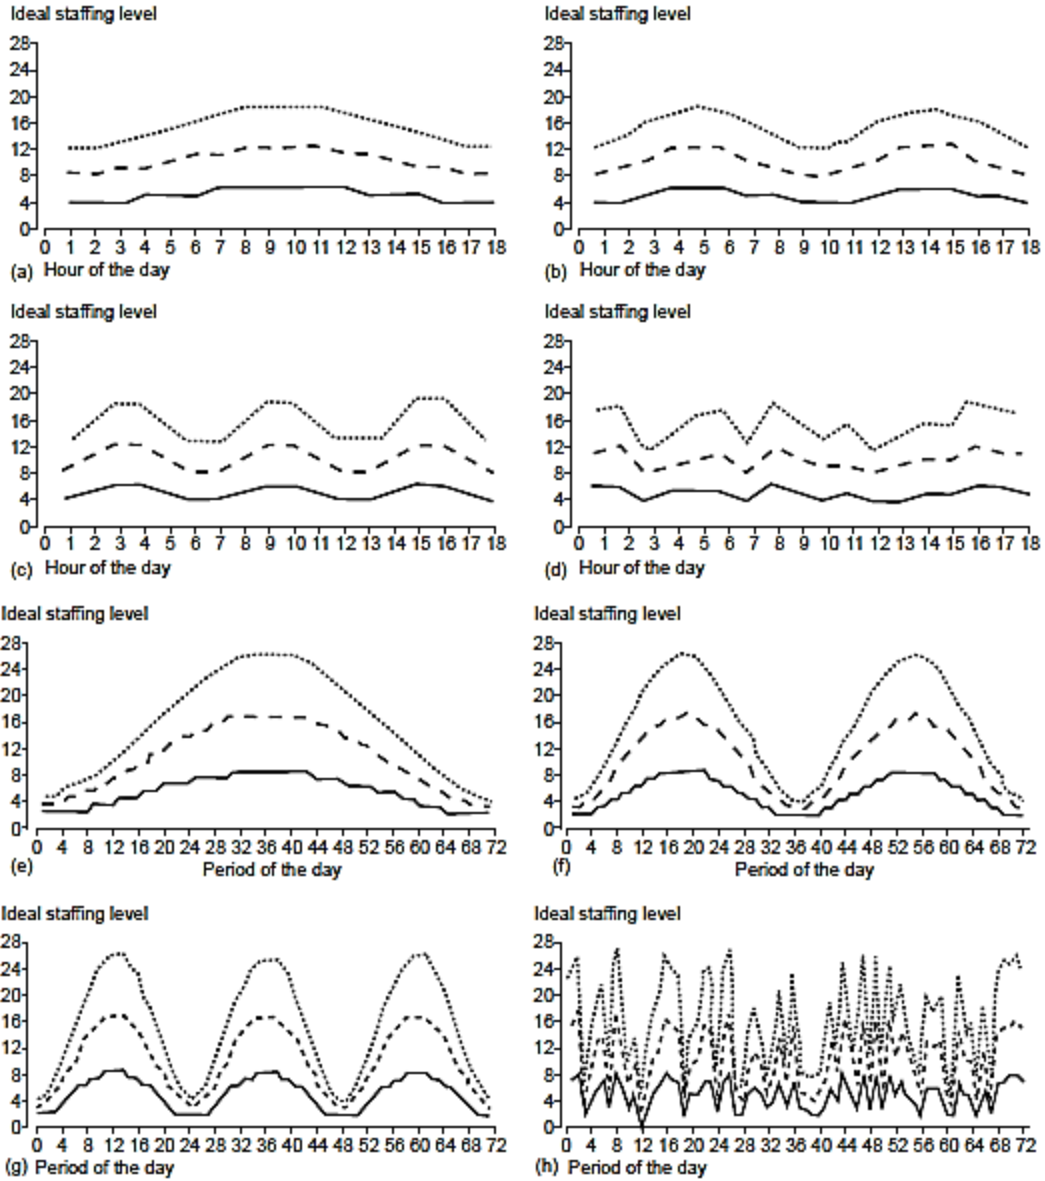
Figure 2. Examples of ideal staffing level patterns in problem set 2. (Unimodal (a,e), bimodal (b,f), trimodal (c,g) and random (d,h) patterns. Patterns (a)-(d) have 60-minute planning intervals, patterns (e)-(h) have 15-minute planning intervals)