

Shift Scheduling in Services When Employees Have Limited Availability:

An L. P. Approach

Gary M. Thompson

David Eccles School of Business

University of Utah

1655 East Campus Center Drive

Salt Lake City, UT 84112

Final version published in: *Journal of Operations Management* Vol. 9, No. 3, Aug. 1990

The author wishes to recognize the insightful comments and suggestions provided by the anonymous reviewers of earlier versions of this paper.

ABSTRACT

This paper compares two linear programming (LP) models for shift scheduling in services where homogeneously-skilled employees are available at limited times.

Although both models are based on set covering approaches, one explicitly matches employees to shifts, while the other imposes this matching implicitly. Each model is used in three forms—one with complete, another with very limited meal break placement flexibility, and a third without meal breaks—to provide initial schedules to a completion/improvement heuristic. The term *completion/improvement heuristic* is used to describe a construction/improvement heuristic operating on a starting schedule.

On 80 test problems varying widely in scheduling flexibility, employee staffing requirements, and employee availability characteristics, all six LP-based procedures generated lower cost schedules than a comparison from-scratch construction/improvement heuristic. This heuristic, which perpetually maintains an explicit matching of employees to shifts, consists of three phases which add, drop, and modify shifts.

In terms of schedule cost, schedule generation time, and model size, the procedures based on the implicit model performed better, as a group, than those based on the explicit model. The LP model with complete break placement flexibility and implicitly matching employees to shifts generated schedules costing 6.7% less than those developed by the from-scratch heuristic.

**Shift Scheduling in Services When Employees Have Limited Availability:
An L. P. Approach**

INTRODUCTION

Services are an exceedingly important component of the modern economy. Thus, in addition to being important at the micro level, good management of services is also important at the macro level. Effective service management has, as one task, the provision of enough personnel to deliver the service. To determine appropriate employee skill levels, a manager must have a good understanding of what customers desire from the service. Since these expectations may vary both across customers and temporally, determining the right type of personnel and ensuring they are present in the necessary quantities are service delivery problems.

Efficient service management requires that just enough appropriately skilled personnel be present to provide the desired level of service. Having more than enough results in excess labor costs, because services cannot be inventoried. Having too few employees may result in poor, customer-annoying service, and risks the loss of future sales. Temporal variation of customer demand compounds the difficulty of providing the necessary number of personnel. The process of matching the numbers and skills of employees actually working at a given time, to the numbers and skills of employees needed at that time in order to provide the desired level of customer support, is commonly known as service labor scheduling.

Many services, especially those utilizing part-time employees, are faced with a pool of staff who are only available for work at certain times during a day or week.

Unfortunately for managers, most labor scheduling research has implicitly assumed employees are continuously available for work, a very different scheduling situation. In addition, much scheduling research has contained restrictive assumptions defining allowable shifts. Because the number of shifts grows combinatorially, common assumptions have been: (1) planning periods of 30 minutes or longer; and (2) small break windows (the allowable periods in shifts when breaks may be taken). In specific instances these restrictions may be valid, although they substantially reduce the scheduling flexibility possessed by many managers. This paper incorporates real-world flexibility by using 15-minute planning intervals and large break windows in some of its test problems.

Two linear programming (LP) models for shift scheduling using homogeneously-skilled employees having limited availability, both utilizing set covering approaches, are compared. The models differ in that the matching of employees to shifts is imposed explicitly in one and implicitly in the other. Each of the models is used in three forms—two with, the third without meal breaks—to provide starting schedules to a completion/improvement heuristic, and are compared to the same heuristic started from scratch. The term *completion/improvement heuristic* refers to a construction/improvement heuristic when operating on a starting schedule.

The structure of this paper is as follows: relevant literature is briefly reviewed; the explicit and implicit models and the from-scratch construction/improvement heuristic are described; the test environment is detailed; model application is described; the results are presented and discussed; and a brief summary concludes the paper.

REVIEW OF RELEVANT LITERATURE

Because of the difficulty in obtaining optimal solutions to problems typical of those occurring in service organizations, heuristic procedures have played the dominant role in service labor scheduling research. The heuristics may be categorized as: (1) those creating schedules from scratch; and (2) those using starting schedules obtained by some other means.

Glover, McMillan, and Glover (1984); Loucks (1986); and Thompson (1990) report the development of from-scratch construction/improvement heuristics incorporating restricted employee availability, each of which uses explicit employee-to-shift assignments. Glover et al. (1984) allow seven types of moves when generating weekly schedules: (1) adding a shift; (2) dropping a shift; (3) modifying a shift; (4) exchanging one shift for another; (5) trading shifts between two people on a given day; (6) moving a shift from one day to another; and (7) cross-trading shifts between two different people on two different days. Loucks (1986) took a somewhat unusual tack to constructing a schedule: the task in a distinct time period which had the fewest number of available people qualified to perform it was identified, and then an employee was scheduled to perform the task at the specified time. Thompson (1990) evaluated an extensive set of criteria for adding shifts to and dropping shifts from a three-phase construction improvement heuristic. To expand the search range arising from an explicit matching of employees to shifts, he allowed three-way shift assignment swaps during schedule construction and improvement. His heuristic is described more fully in a later section.

Although no reported heuristic incorporating restricted employee availability has utilized starting solutions, this approach has been commonly applied when employees are

assumed to be continuously available. Henderson and Berry (1976) and Morris and Showalter (1983) demonstrated the value of using, as initial schedules for heuristics, solutions provided by LP relaxations of models based on Dantzig's (1954) shift scheduling formulation. Henderson and Berry (1976) evaluated methods of selecting working subsets of shifts in a very flexible scheduling environment. Only shifts in the working subset were used to construct and improve the schedule. In LPVE, the best performing of the heuristics they evaluated, any fractional shift decision variable values in the LP relaxation solution were rounded up. The schedule was then improved, if possible, by eliminating two shifts concurrently with the addition of a third. Morris and Showalter obtained initial schedules by rounding down fractional decision variable values, and evaluated several heuristic rules for improving the schedule.

Keith (1979), who also used LP to generate initial schedules, rounded any fractional shift variable values to the nearest integer. Although he eliminated break placement flexibility in the LP model to reduce its size and accelerate its solution, complete break placement flexibility was allowed, and taken advantage of, during subsequent schedule improvement. Other researchers (Segal (1974), Moondra (1976), Bailey (1985), Love and Hoey (1990)) have facilitated solving LP models by approximating shifts with breaks by shifts without breaks. Of these, only Segal (1974) presents a procedure (a separate network model) for accommodating breaks. A contrasting approach has been provided by Bechtold and Jacobs (1990): an LP model in which breaks are implicitly matched to shifts. When ample break scheduling flexibility existed, this approach enabled a substantial reduction in model size, and resulted in their

obtaining optimal solutions to what would otherwise have been quite large shift scheduling problems.

Love and Hoey (1990) presented two network flow models to be used sequentially for scheduling employees having limited availability. The first model, which assumed all employees were continuously available, scheduled shifts with the objective of minimizing surplus staffing. The second model assigned employees to the shifts scheduled by the first model. It had as multiple objectives the minimization of: (1) task understaffing; (2) inefficient, undesirable, and impossible shift assignments; and (3) deviations from the desired number of working days for each employee. The disadvantage of Love and Hoey's two-phase approach is that there is no guarantee that the schedule generated by the first model is availability-feasible. *Availability-feasibility* implies that all shifts in a schedule can be staffed despite the limited availability of employees.

This paper extends published research in labor scheduling by using, when employees are limited in their availability, linear programming to provide initial schedules to a schedule completion/improvement heuristic. Importantly, the initial schedules are guaranteed to be availability-feasible by the nature of the LP models.

DESCRIPTION OF THE SHIFT SCHEDULING PROCEDURES

Within this section the two LP models are introduced and the from-scratch heuristic is described.

Explicit Employee-to-Shift Matching Model (EMLP)

Thompson (1990) modified Dantzig's (1954) integer programming shift scheduling model for employees having limited availability. The model (EMLP), which explicitly matches employees to shifts, is:

$$\text{Minimize } Z = \sum_{p \in P} (cu_p \cdot u_p + co_p \cdot o_p)$$

subject to

$$\sum_{e \in E} \sum_{i \in SS_e} a_{ip} \cdot x_{ei} + u_p - o_p = r_p \quad \text{for } p \in P$$

$$\sum_{i \in SS_e} x_{ei} \leq 1 \quad \text{for } e \in E$$

$$x_{ei} = 0, 1 \quad \text{for } e \in E; i \in SS_e$$

$$u_p, o_p \geq 0 \quad \text{for } p \in P$$

where

- p is an index for periods
- P = the set of planning periods to be scheduled
- e is an index for employees
- E = the set of available employees
- i is an index for shifts
- cu_p = understaffing cost per employee in period p
- co_p = overstaffing cost per employee in period p
- u_p = shortage of employees in period p
- o_p = surplus of employees in period p

SS_e	= the set of shifts for which employee e is available to work
x_{ei}	= $\begin{cases} 1, & \text{if employee } e \text{ is assigned to shift } i \\ 0, & \text{otherwise} \end{cases}$
a_{ip}	= $\begin{cases} 1, & \text{if } p \text{ is working (i.e., non - break) period for shift } i \\ 0, & \text{otherwise} \end{cases}$
r_p	= number of employees needed to provide the desired level of customer service in period p

The objective, (1), is to schedule the employees at a minimum possible cost. Costs are incurred for having excess as well as insufficient numbers of employees on-hand. Constraint set (2) requires that for each period, the number of employees on-hand, less any surplus and plus any shortage, equals the number of employees desired. It is necessary to allow understaffing since limited employee availability may preclude satisfying staffing requirements in some periods. The under- and overstaffing variables in the objective result in EMLP having the form of a single-priority, weighted-deviational-variable, mixed-integer goal programming model. As EMLP contains no preemptive priorities, its LP relaxation may be solved using commonly available LP software.

Constraint set (3) restricts each employee to working no more than one shift. Constraint set (4) imposes the binary integrality of the employee-shift variables, which results in the model having an explicit match of employees to shifts. Constraint set (5) specifies the non-negativity of the employee shortage and surplus variables.

In those test problems in which any employees had identical availability, a size reduction procedure was applied to EMLP: (1) shift variables were defined for only one of the identically available employees; (2) the right hand side of the appropriate constraint from set (3) was set to the number of employees having identical availability;

and (3) the appropriate shift variables in constraint set (4) were changed to general integer.

Implicit Employee-to-Shift Matching Model (IMLP)

The model implicitly matching employees to shifts (IMLP) was developed from the insight that if the availability of two or more employees overlap, the intersection of the sets of shifts they can work is unlikely to be empty. In this event, the need for unique shift variables for every employee is eliminated. To facilitate the modelling process, regions are defined to be sets of shifts which can be staffed by the same subset of the available employees. These regions are disjoint and their union contains all shifts, since one of the regions will contain those shifts which cannot be staffed by any employee. In the model presented below, employees, based on their availability, are first mapped onto the regions. In turn, the regions are mapped onto the set of staffable shifts. Define:

- v_g = the earliest starting time for any shift in region g
- w_g = the latest finishing time for any shift in region g
- f_e = the first period in which employee e is available for work
- l_e = the last period in which employee e is available for work

An employee, e , can work any region, g , providing that: (1) his/her earliest availability does not exceed the earliest starting time of any shift in the region ($f_e \leq v_g$); and (2) his/her latest availability is not less than the latest finishing time of any shift in the region ($l_e \geq w_g$). An algorithm for identifying regions and their characteristics is provided in Appendix A.

Using the implicit matching of employees to shifts, the IMLP model is:

subject to

$$\sum_{n \in N} a_{np} \cdot s_n + u_p - o_p = r_p \quad \text{for } p \in P$$

$$\sum y_{eg} \leq 1 \quad \text{for } e \in E$$

$$\{g \in G \mid f_e \leq v_g, w_g \leq 1_e\}$$

$$\sum k_{ng} \cdot s_n - \sum y_{eg} = 0 \quad \text{for } g \in G$$

$$n \in N \quad \{e \in E \mid f_e \leq v_g, w_g \leq 1_e\}$$

$$y_{eg} = 0, 1 \quad \text{for } e \in E; \{g \in G \mid f_e \leq v_g, w_g \leq 1_e\}$$

$$s_n \geq 0 \text{ and integer} \quad \text{for } n \in N$$

where

n	is an index for shifts
N	= set of distinct shifts
s_n	= the number of employees working shift n
a_{np}	= $\begin{cases} 1, & \text{if period } p \text{ is a working period of shift } n \\ 0, & \text{otherwise} \end{cases}$
g	is an index for regions
G	= set of regions
y_{eg}	= $\begin{cases} 1, & \text{if employee } e \text{ works some shift falling in region } g \\ 0, & \text{otherwise} \end{cases}$

$$k_{ng} = \begin{cases} 1, & \text{if shift } n \text{ falls in region } g \\ 0, & \text{otherwise} \end{cases}$$

with $p, P, u_p, o_p, r_p, f_e, l_e, v_g, w_g, e,$ and E as defined earlier

As before, the objective (1) is to minimize the sum of under- and overstaffing costs. Constraint set (6) ensures that, for each period, the number of employees working, adjusted for shortages and surpluses, equals the number desired. Constraint set (7) limits each employee to working a shift from only one region. Constraint set (8) equates, for each region, the number of employees staffing the region to the number of scheduled shifts from the region. It is this constraint set that captures the implicit matching of employees to shifts. Appendix B presents a simple scheduling problem, illustrating a graphical representation of regions and showing constraint sets (7) and (8).

A model size reduction process, valid if any employees have the same availability, was applied whenever possible in this paper: (1) separate employee-region variables were defined for only one of the employees having identical availability; (2) the right hand side value of the appropriate constraint from set (7) was fixed at the number of identically available employees; and (3) the appropriate employee-region variables in constraint set (9) changed from binary to general integer.

A From-Scratch Shift Scheduling Heuristic (FSH)

The three-phase construction/improvement heuristic developed by Thompson (1990) is utilized as a control technique to develop schedules from scratch. (As will be described later, this heuristic is also used for completing and improving initial schedules obtained from solutions to the LP relaxations of the EMLP and IMLP models.) This

heuristic incorporates restricted employee availability by perpetually maintaining an explicit match of employees to shifts.

In phase 1 of the heuristic, shifts are added to the schedule, while others may be dropped from the schedule in phase 2. Phase 3 improves the schedule by moving breaks; displacing, lengthening, and shortening shifts; and through other, more complex changes to the schedule. Hence, the actions in the three phases of this heuristic are comparable to the moves employed by in the construction/improvement heuristic of Glover et al. (1985).

The criteria used for adding and dropping shifts were those found by Thompson (1990) to generate the best schedules in his heuristic. During schedule construction, the shift with the tightest availability index is added to the schedule. This index is defined as the average, across all working periods in the shift, of the net employee requirements, divided by the number of employees available to work the period and unassigned to any shift; with the term *net employee requirement* defined as the number of employees required in a period, less the number currently scheduled to work in that period. Note that the availability index is somewhat similar to the selection criterion used by Loucks (1986). Ties are broken first by choosing the shift resulting in the smoothest net staffing requirements, and second by selecting the longest shift. Smoothness is defined as:

$$\text{schedule smoothness} = \sum_{p=1}^{p+1} |n_p \cdot n_{p-1}|$$

where

P	= number of periods to be scheduled
n_p	= current new employee requirements in period p
$ x $	= absolute value of x
n_o	= $n_{p+1} = 0$

During phase 2, the criterion of maximizing the schedule cost reduction per working period of a shift (or minimizing the per-period cost increase) is used to identify shifts to be dropped from the schedule. Ties are broken by choosing the shift which, if dropped, yields the greatest reduction (least increase) in schedule cost. If the schedule cost would be lowered by eliminating the selected shift, the drop is made; otherwise, a second shift is also dropped, and if beneficial, new shifts added to the schedule. In this way, schedule improvements, such as those employed by Henderson and Berry (1976) in their LPVE heuristic, will occur.

Phase 3 of the heuristic is comprised of seven separate procedures. Four—those which move breaks, and lengthen, shorten, or offset shifts—are used in a rotating sequence until no further improvement can be made. Through manipulations of the under- and overstaffing costs per period, the remaining three processes attempt: (1) to make quasi-simultaneous shift extensions and contractions; and (2) to add or (3) drop a shift while making accommodating adjustments to other shifts. In the improvement activities which lengthen or shorten shifts, meal break duration may also be adjusted. Neither availability-feasibility nor shift-defining restrictions are violated at any time during schedule construction, completion, or improvement.

TEST ENVIRONMENT

In all test problems: (1) a 20-hour operating day was used; (2) employee requirements were assumed known with certainty for all planning periods; (3) under- and overstaffing costs per employee-period were assumed to be constant across the operating day, and equal to 1 and 4, respectively; (4) every employee was assumed to be able to

work a shift of any duration for which he/she was available; (5) rest breaks (breaks shorter than 30 minutes) were not scheduled; (6) two shift types were used—shifts of four to six working hours received a half-hour meal break, while shifts of six to eight working hours received an hour-long meal break; and (7) split shifts (shifts with meal breaks longer than one hour) were not considered.

Five factors were varied in developing the test data, yielding a total of 80 test problems: (1) the employee requirements pattern (five patterns); (2) the mean employee requirements (two levels); (3) the average availability of employees (two levels); (4) coverage—the ratio of the maximum total hours that could be worked by all employees to the total labor hours required to satisfy expected customer demand (two levels); and (5) the degree of scheduling flexibility (two levels).

Uniform, unimodal, bimodal, trimodal, and pattern-less (i.e., random) employee requirements curves that are representative of the range of customer arrival patterns occurring in services were utilized. With the exception of uniform demand, all patterns have coefficients of variation of approximately 0.3 and are illustrated in Figure 1 for each of the two levels of mean employee requirements.

Low and high mean employee requirements were chosen as 5 and 10 employees, respectively. At the low and high levels of employee availability, employees were available for an average of 7 and 14 hours, respectively, and divided into three approximately equal groups. Individuals in these groups were available for work during 5, 7, and 9 (10, 14, 18) hours of the operating day at the low (high) level of employee availability. The low and high coverage values were 1.25 and 2.50, respectively. Clearly,

values closer to one result in less flexibility in developing schedules (a tighter labor supply) than do higher values.

Insert Figure 1 Here

Insert Table 1 Here

Table 1 gives the work force sizes under the different levels of the various test factors. The problems range from a low of 16 (when mean requirements are for five employees per period, coverage is low, and average employee availability is high) to a high of 81 available employees (when coverage is high, average employee availability is low, and ten employees are required each period, on average).

Table 2 summarizes the restrictions operative under the two levels of scheduling flexibility. These restrictions are representative of those existing in many service organizations. At the high and low levels of scheduling flexibility, respectively: (1) the operating day was broken into eighty 15-minute, and forty 30-minute planning intervals; and (2) ignoring availability- feasibility, the scheduling restrictions result in maximums of 495 and 6588 distinct shifts that could be staffed.

Insert Table 2 Here

APPLICATION OF THE EMLP AND IMLP MODELS

Both the EMLP and IMLP models are evaluated in three forms: (1) AB, which incorporates complete break placement flexibility; (2) SB, which allows a sole position for a break; and (3) NB, in which breaks are not scheduled. Eliminating break placement flexibility has the effect of reducing the size of the resultant LP model, and increasing the likelihood that it can be solved. As was done by Keith (1979), *full break placement flexibility is incorporated during the succeeding heuristic schedule completion and improvement.*

The following sequence of activities was performed for all of the LP-based scheduling procedures: (1) the LP relaxation of the appropriate model was generated and solved; (2) any fractional shift variable values in the LP solution were rounded down; (3) meal breaks were inserted into shifts, if necessary; (4) employees were assigned to shifts, if required; and (5) the resultant schedule was completed and improved, if possible. Each of these actions is described more fully below.

Generating and Solving the Models

Two issues relating to the generation of the models merit mention: the region identification process used for the IMLP models and the actions taken to keep the model size manageable. In generating the IMLP models, a FORTRAN subroutine based on the

algorithm in Appendix A was used to identify regions and their characteristics. Although it may be possible to develop a faster procedure for identifying regions, the algorithm provided in Appendix A works quite effectively.

Obviously, attempts could be made to solve the EMLP and IMLP models to optimality. Their size, however, makes such an endeavor exceedingly difficult. Hence, we resort to solving the LP relaxations of the models. As the degree of scheduling flexibility increases, so does the size of the models. If 15-minute planning intervals had been used in the high flexibility scheduling environment in this paper, all of the EMLP and IMLP models would have been too large to solve using SAS-OR on an IBM 3090 computer. Because of this, the planning intervals in the LP models were raised from 15 to 30 minutes. *Complete scheduling flexibility was restored during the ensuing heuristic completion and improvement.* Hence, while the complete set of shifts is usable during heuristic schedule improvement, only a working subset of shifts is considered when solving the LP relaxations of the EMLP and IMLP models.

In some of the test problems, the resultant models were still too large to be solved (particularly the EMLP-AB models). In these cases, there would be no initial schedule and so the construction/improvement heuristic would be used to generate schedules from scratch.

Rounding Down Fractional Shift Variable Values

The literature has contained examples of rounding up (Henderson and Berry (1976)), rounding down (Morris and Showalter (1983)), and rounding fractional shift variable values to the nearest integer (Keith (1979)). When scheduling employees having

limited availability, fractional shift values must be rounded down in order to preserve the availability-feasibility of the schedule.

Inserting Meal Breaks into Shifts

It is only necessary to insert meal breaks into shifts from the NB versions of the EMLP and IMLP models, since it is only the NB model versions which fail to schedule breaks. The meal break insertion heuristic has the objectives of: (1) minimizing the schedule cost, and (2) making the resultant net employee requirements as smooth as possible, as measured using (11). During schedule improvement, smooth net employee requirements facilitate the extension of shifts to cover periods of understaffing and the contraction of shifts from periods of overstaffing.

Matching Employees to Shifts

Since the heuristic improvement procedure requires an explicit matching of employees to shifts, whenever the IMLP model is used employees must be assigned to shifts prior to commencing the heuristic completion and improvement activities. In matching employees to shifts: (1) shifts are ordered from early to late starts, with ties broken randomly; and (2) employees are allocated to the ordered shifts by selecting from the subset of assignable employees, the person whose last period of availability most closely matches the last period in a shift. Note that it is not important how early an assignable employee is available, since all shifts with earlier starts have been allocated. However, some as yet unallocated shifts may finish later than the shift currently being assigned, so employees available later should be conserved.

Heuristically Completing and Improving the Schedule

In all cases, heuristic completion and improvement of the initial schedule is undertaken using the FSH heuristic previously described. Completion and improvement are attempted for the following reasons: (1) at the high level of scheduling flexibility, all the LP models used 30- minute planning intervals (using the appropriate 15-minute intervals may enable schedule improvement); (2) break placement flexibility is not incorporated in the SB and NB versions of the models (utilizing this flexibility may facilitate schedule improvement); (3) the fractional shift values in the LP relaxation solutions that are rounded down represent shifts lost (and thus offer the potential for adding new shifts or modifying those remaining); (4) some of the models were too large to be solvable (in these cases, the heuristic was used to generate schedules from scratch).

PRESENTATION AND DISCUSSION OF RESULTS

Within this section, schedule cost, computational time, and model size results are presented. As well, an extensive discussion is offered which focuses on: (1) observations based on and (2) suggestions for improving the performance of the NB versions of the LP models; (3) overall performance of the scheduling procedures; and (4) extensions and modifications to IMLP.

Schedule Costs

Detailed performance results are presented in Table 3, while Table 4 presents summary results. FSH, the construction/improvement heuristic started from scratch, yielded higher cost schedules than all of the LP-based procedures. These results are

consistent with the findings of Henderson and Berry (1976) and Morris and Showalter (1983). Compared to a from-scratch heuristic, Morris and Showalter (1983) found that their LP-based heuristic reduced average schedule costs by approximately 1%. In contrast, EMLP-AB, the poorest performing of the LP-based procedures, reduced mean schedule costs by 3.0% compared to the from-scratch FSH heuristic, while IMLP-AB, the best performing LP-based procedure, reduced schedule costs by 6.7%. The better relative performance of the LP-based procedures may be a result of the increased complexity of the limited employee availability environment, and a consequential decrease in the effectiveness of heuristic schedule construction.

Computational Times

Schedule generation time is defined as the total time required to perform all activities associated with developing a schedule. For IMLP and EMLP it thus includes the time required to: (1) develop and solve the LP models (including region identification in IMLP); (2) obtain a starting schedule by interpreting the model solution; and (3) heuristically complete and improve the initial schedule.

In terms of relative computational times, the LP-based procedures required between 4.5 and 6.5 times the CPU seconds used by FSH to generate schedules, on average. In absolute terms, schedule generation times ranged from 7.3 to 10.6 seconds on an IBM 3090 computer, on average. As such, developing schedules using LP relaxations of the EMLP and IMLP models cannot be viewed as computationally burdensome.

Model Size

A comparison of the size of the AB and NB versions of the EMLP and IMLP models is presented in Table 5. IMLP-AB (IMLP-NB) averaged 48.4% (90.3%) of EMLP-AB (EMLP- NB) in terms of A matrix size and had only 14.4% (16.8%) of the A-matrix non-zeros. Generally, the relative size advantage of IMLP improved as the size of EMLP increased, but particularly as the break placement flexibility increased. As the range of break timing options increased from none to very flexible, the size of IMLP approximately doubled, but EMLP approximately tripled in size. This size advantage occurred with increased scheduling flexibility because the implicit nature of IMLP enabled a more compact mapping of employee availability onto the shift alternatives than EMLP's explicit matching of employees to shifts.

Insert Table 3 Here

Insert Table 4 Here

The main drawback of EMLP is the large number of employee-shift variables required, particularly in those cases where scheduling flexibility is high. Indeed, even though 30-minute planning intervals were used in all the LP models at the high level of scheduling flexibility, EMLP-AB (EMLP-SB and EMLP-NB) was (were) too large to be solvable in 41 (10) of the 80 test problems. In contrast, all versions of IMLP were

solvable in all of the test problems. This shortcoming of EMLP is a practical concern, since conditions of high scheduling flexibility (including employees being available for work over long periods and the use of small planning intervals, numerous allowable shift lengths, and large break windows) occur in many service organizations.

In essence, increasing the planning interval duration in EMLP and IMLP has the effect of restricting the set of shifts that can be utilized. In other words, only a working subset of shifts is used during the solution of the LP relaxations of these models.

Although this action was effective (compare, from Table 3, the schedule costs for the LP-based procedures to those for FSH at the high level of scheduling flexibility), other means are available for ensuring the models are of a solvable size. One possibility is to use alternate methods of selecting working subsets of shifts. A second option would utilize implicit modeling of breaks (along the lines of the work of Bechtold and Jacobs (1990)). Specifically for the environment of non-continuously available employees, the relative and absolute effectiveness of these, or other mechanisms, merits investigation.

Insert Table 5 Here

Observations Based on the Performance of EMLP-NB and IMLP-NB

Examining the summary results in Table 4, it is noteworthy that the initial schedules from the NB model versions are of lower cost than the initial schedules from the AB or SB versions of the same model. The lower costs occurred because fewer partial shifts (in the solutions to the LP relaxations of the NB model versions) were lost when fractional shift values were rounded down. These results suggest that EMLP-NB and

IMLP-NB at least partially preserve the network structure of the unlimited-employee-availability, no-break shift scheduling problem (see Segal (1974)).

Even though IMLP-NB had the lowest cost initial schedules, the best performing procedure used IMLP-AB, indicating that shifts without breaks fail to accurately represent shifts with break placement flexibility. (That EMLP-NB outperformed EMLP-AB may be dismissed, because EMLP-AB was not solvable in 41 of the 80 test problems). This conclusion has broad implications since: (1) various scheduling techniques presented in the literature failed to schedule meal breaks (Baker, Crabill, and Magazine (1973); Bailey (1985); Loucks (1986); Love and Hoey (1990)); and (2) service facilities where employees are not given meal breaks (for shifts over six hours, anyway) are likely to be uncommon.

Possibilities for Improving the Performance of EMLP-NB and IMLP-NB

Various options exist for attempting to improve the performance of EMLP-NB and IMLP-NB. All strive to alleviate the effects of a lack of breaks in these models. First, employee requirements may be increased, in those periods where breaks will likely be necessary, by the expected number of break periods required. True employee requirements would be reinstated during heuristic break insertion and schedule improvement. Second, some employees may be prohibited from being scheduled by the LP models. These employees could then be used during schedule improvement to cover the required breaks. Third, restrictions limiting the number of shifts changeovers (shifts starting and finishing at the same time) may be incorporated in the models. Such restrictions would be helpful as EMLP-NB and IMLP-NB have a tendency to schedule

many identical shifts, particularly when employee requirements are uniform. This commonality of shifts results in breaks being scheduled in a short time stretch and fetters the schedule improvement heuristic in extending shifts to cover breaks. The evaluation of these procedures, or others, is left to future research.

Overall Performance

As a group, the performance of the IMLP-based procedures was distinctly superior to that of the EMLP-based procedures: (1) all of the IMLP models were solvable, while 1/8 of the EMLP- SB and EMLP-NB, and more than 1/2 of the EMLP- AB models were not; (2) on average, IMLP- AB and IMLP-NB generated the lowest and next lowest cost schedules, respectively; (3) schedule generation times were lower for the IMLP procedures than for the EMLP procedures (the low time for EMLP-AB is because of the large number of LP models that were too large for solution to be attempted); and (4) on average, the IMLP models were smaller. Given its superiority, IMLP will be the focus of the remainder of the paper.

Extensions and Modifications to IMLP

The scheduling test data did not contain split availability within days for any employee. In reality, split availabilities likely occur frequently. Fortunately, IMLP is easily adapted to handle this. Employee-region variables can be defined for each availability stretch. Summing the employee-region variables over all availability stretches in constraint set (7) would ensure that the employee still works at most a single shift.

In the scheduling environment considered, any employee could work a shift of any duration, providing his/her availability would allow it. In some situations, different classes of employees may exist (such as full-time and part-time employees) which work shifts of non-identical durations. A simple way of dealing with this in IMLP would, for each of the separate employee classes, define separate shifts and use the algorithm in Appendix A to identify the relevant regions.

With modifications, IMLP might be appropriate for tour scheduling. A weekly tour scheduling model could be constructed as a composite of the separate daily shift scheduling problems. Additional restrictions would be necessary to impose weekly constraints. For example, a limit could be placed on the number of days an employee could work. Since the match of employees to shifts is implicit, directly imposing restrictions on the maximum hours per week for employees would be difficult. As such, this extension is left to future research. Alternately, the IMLP model could be used in a disaggregation approach to tour scheduling, whereby weekly schedules are built up from the solutions to the daily shift scheduling problems. Such approaches make more sense in environments like the one studied, rather than in those where employees are continuously available, since employee availability frequently varies across days in service organizations.

SUMMARY

This paper has made several contributions to the field of labor scheduling: (1) linear programming has been used for the first time to generate initial shift schedules for services where employees are available only at limited times; (2) a new model, IMLP,

which implicitly matches employees to shifts has been introduced; and (3) both EMLP, a model which explicitly matches employees to shifts, and IMLP generate schedules for which all shifts can be assigned to individual employees, despite the limited availability of employees. The results on a wide range of test problems illustrate the superiority—in terms of schedule cost, schedule generation time, and model size—of procedures based on the IMLP model over those based on the EMLP model. Overall, the procedure using the IMLP model version which incorporates total break placement flexibility performed best, generating schedules 6.7% better than those developed by a from-scratch heuristic.

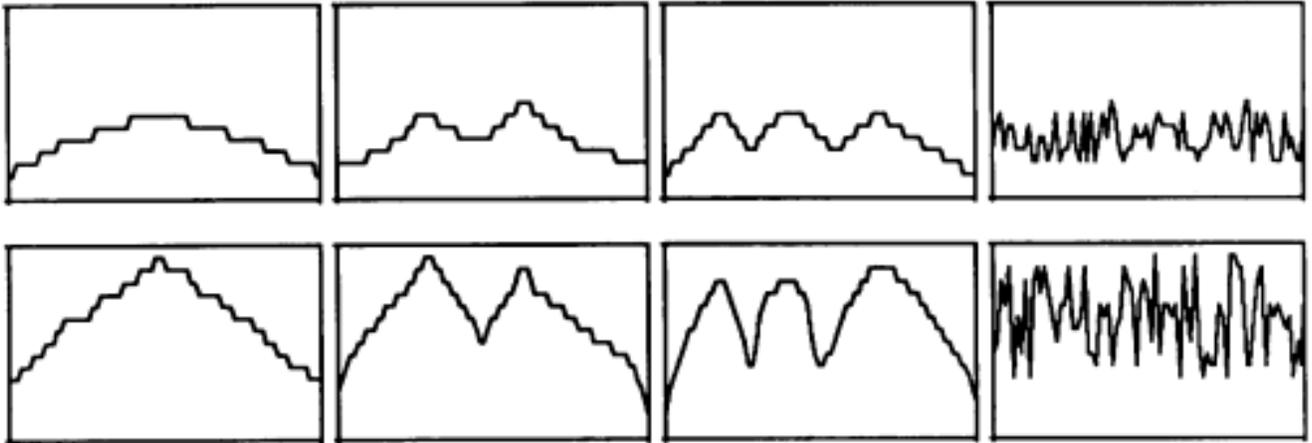


Figure 1. Employee requirements patterns.

(The scale of all the graphs are identical, ranging from 0 to 16 employees.)

Table 1. Workforce sizes in the test problems.

		Average Employee Availability			
		7 hours		14 hours	
Coverage*		Low	High	Low	High
Average per-period Employee	5	20	41	16	31
Requirements	10	41	81	31	63

*Coverage is defined as the ratio of the maximum total hours that can be worked by all employees to the total labor hours required to provide the desired level of customer service.

Table 2. Restrictions operative under the two levels of scheduling flexibility.

Flexibility Level:		Low	High
Planning interval length (minutes)		30	15
Number of planning intervals in a 20 hr day		40	80
Shift working length ^a increments (minutes)		30	15
Shift start-time increments (minutes)		30	15
Break start-time increments (minutes)		30	15
Minimum, maximum work stretches before and after a break (minutes)	half-hour breaks	120,180	90,210
	hour-long breaks	180,240	150,270
Number of distinct shifts (with a 20-hr. day) ^b		495	6588

^aWorking length is the duration of the shift excluding any breaks.

^bThe number of unique shifts that could be scheduled if all employees were continuously available for work.

Table 3. Detailed information on schedule costs and generation times

Scheduling flexibility (L = low, H = high)					Scheduling Procedure						
Mean requirements (L = 5, H = 10 employees)											
Average employee availability (L = 7, H = 14 hours)											
Coverage (L = 1.25, H = 2.50) ^a											
C = Average schedule cost (per period)											
T = Average total schedule generation time (seconds) ^b											
					EMLP			IMLP			FSH
					AB	SB	NB	AB	SB	NB	
L	L	L	L	C	2.110 ^c	2.195	2.175	2.085	2.175	2.135	2.135
L	L	L	H	C	0.605	0.615	0.615	0.625	0.570	0.650	0.605
L	L	H	L	C	0.855	0.980	0.870	0.850	0.975	0.695	1.010
L	L	H	H	C	0.215	0.315	0.300	0.280	0.300	0.235	0.215
L	H	L	L	C	2.610	2.495	2.715	2.565	2.485	2.780	2.785
L	H	L	H	C	0.260	0.310	0.300	0.300	0.275	0.350	0.290
L	H	H	L	C	1.330	1.040	1.010	0.855	1.055	0.970	1.330
L	H	H	H	C	0.225	0.255	0.255	0.190	0.155	0.275	0.225
H	L	L	L	C	1.845	1.855	1.848	1.868	1.785	1.833	1.845
H	L	L	H	C	0.590	0.623	0.643	0.610	0.638	0.585	0.595
H	L	H	L	C	1.108	0.933	0.778	0.920	0.948	0.793	1.108
H	L	H	H	C	0.260	0.285	0.295	0.278	0.268	0.350	0.260
H	H	L	L	C	2.360	2.378	2.320	2.435	2.415	2.295	2.485
H	H	L	H	C	0.433	0.393	0.508	0.438	0.488	0.458	0.433
H	H	H	L	C	1.268	1.158	1.095	1.153	1.175	1.075	1.268
H	H	H	H	C	0.363	0.363	0.363	0.360	0.390	0.415	0.363
L	L	L	L	T	4.94 ^c	3.49	3.05	4.72	3.67	2.91	0.67
L	L	L	H	T	8.25	5.80	4.92	8.28	7.37	5.92	0.79
L	L	H	L	T	12.71	9.73	4.21	5.36	4.29	3.83	0.57
L	L	H	H	T	1.36	13.58	13.43	7.90	6.51	5.87	0.63
L	H	L	L	T	9.25	6.69	6.44	9.72	7.45	7.58	2.66
L	H	L	H	T	13.35	9.01	8.13	14.62	13.21	10.75	1.49
L	H	H	L	T	2.17	15.25	15.33	7.21	5.94	5.60	1.47
L	H	H	H	T	1.91	1.55	1.73	11.37	9.94	9.55	0.89
H	L	L	L	T	9.65	3.86	3.36	16.52	9.49	8.05	0.97
H	L	L	H	T	16.63	6.87	6.12	12.69	8.20	7.40	1.65
H	L	H	L	T	1.89	10.34	9.88	8.49	4.61	4.26	0.92
H	L	H	H	T	3.36	15.03	15.09	11.68	7.47	7.14	1.88
H	H	L	L	T	17.70	8.80	7.67	13.43	9.11	8.86	3.23
H	H	L	H	T	4.37	12.13	10.95	21.28	15.37	15.74	3.51
H	H	H	L	T	3.61	16.09	15.63	10.78	6.65	6.02	2.21
H	H	H	H	T	4.58	3.29	3.24	15.83	11.20	12.43	2.50

^aCoverage is defined as the ratio of the maximum total hours that can be worked by all employees to the total labor hours required to provide the desired level of customer service.

^bComputer times are for an IBM 3090, with the FSH heuristic coded in FORTRAN and the LP models solved using SAS-OR.

^cAll values are averages over the five employee requirements patterns.

Table 4. Summary results (all values are averages).

	Scheduling Procedure						FSH
	EMLP			IMLP			
	AB	SB	NB	AB	SB	NB	
Schedule cost (per period)	1.027	1.010	1.004	0.988	1.006	0.993	1.059
Schedule generation time (seconds)	7.23	8.84	8.39	10.62	7.80	7.32	1.63
Per-period schedule cost reduction from solving the LP model*	4.99	10.29	13.23	11.91	13.74	20.24	NA
Per-period schedule cost reduction made by the completion/construction and improvement heuristic*	23.99	18.71	15.77	17.10	15.26	8.77	28.94
The percentage of the total possible schedule cost reduction that resulted from the solution of the LP model*	16.6	34.3	44.1	39.7	45.8	67.5	NA
The percentage of the total possible schedule cost reduction that was made by the completion/construction and improvement heuristic*	80.0	62.4	52.6	57.0	50.9	29.2	96.5

*The average cost of a null schedule (i.e., with no employees scheduled) is 30 units per period. If, after solving the LP model and rounding down any fractional shifts, the cost of the initial schedule is 12 units per period, the LP model would have reduced the schedule cost by 18 units per period. This would represent a schedule cost reduction of 60% of the total possible.

Table 5. A comparison of the size of the LP shift scheduling models on the test problems.

Scheduling Flexibility (L = low, H = high) ^a				LP Model and Version							
Mean Requirements (L = 5, H = 10 employees)											
Average Employee Availability (L = 7, H = 14 hours)											
Coverage (L = 1.25, H = 2.50) ^b											
				EMLP-AB		IMLP-AB		EMLP-NB		IMLP-NB	
				#EL	#NZ	#EL	#NZ	#EL	#NZ	#EL	#NZ
L	L	L	L	47219 ^c	8938	52355	4717	28225	5510	36749	2983
L	L	L	H	114779	17492	155158	6222	66893	10742	123411	4142
L	L	H	L	198315	45373	67035	6848	111997	27981	48599	4432
L	L	H	H	359874	71562	140473	7708	202113	44113	115613	5290
L	H	L	L	103553	16423	129848	6070	60287	10083	100233	3957
L	H	L	H	218450	27373	329008	7659	125845	16804	281622	5354
L	H	H	L	329890	68466	97342	7238	185357	42220	75958	4822
L	H	H	H	607479	105287	259576	8863	340027	64900	229656	6447
H	L	L	L	102075	18797	90934	9563	28225	5510	36749	2983
H	L	L	H	251664	36805	232850	11987	66893	10742	123411	4142
H	L	H	L	414109	95519	113125	13858	111997	27981	48599	4432
H	L	H	H	754278	150694	202623	14716	202113	44113	115613	5290
H	H	L	L	229407	34556	201499	11832	60287	10083	100233	3957
H	H	L	H	504234	57663	443997	13947	125845	16804	281622	5354
H	H	H	L	691224	144188	150802	14248	185357	42220	75958	4822
H	H	H	H	1276109	221765	334376	15873	340027	64900	229656	6447
Average				387666	70056	187563	10084	140093	27794	126480	4678

^aModel sizes in the high flexibility cases are based on 30-minute planning intervals.
^bCoverage is defined as the ratio of the maximum total hours that can be worked by all employees to the total labor hours required to provide the desired level of customer service.
^cAll values are averages over the five employee requirements patterns.

(#EL= number of elements in the A-matrix, #NZ=number of non-zero A-matrix elements)

APPENDIX A

This appendix presents an algorithm for identifying employee regions. To simplify the presentation of the algorithm, it is assumed that each shift may be uniquely identified by its overall duration (including any breaks) and its starting time. It is also assumed that every employee can work any shift for which he/she is available. Begin by defining:

- c = an index for the starting times of shifts,
- d = an index for the overall duration of shifts,
- d_{\max} = the maximum overall duration in periods (including breaks) of any shift,
- d_{\min} = the minimum overall duration in periods (including breaks) of any shift,
- i = a counter,
- $m_{(d,c)}$ = the region to which the shift of d periods in overall duration commencing in period c is assigned,
- $n_{(d,c)}$ = an index number for the shift of d periods in overall duration which commences in period c ,
- P = the number of periods to be scheduled,
- q = an index for the finishing times of shifts.

Using these terms and the definitions presented in the main paper, the algorithm for identifying the regions and their characteristics consists of the following 15 steps:

- (1) Set $g = 0$, $i = 0$, $d = d_{\max} + 1$, and $k_{jg} = 0$ for $j, g = 1, \dots, \infty$. Continue.
- (2) Set $d = d - 1$, $c = 1$, $q = d$, $i = i + 1$, and $n(d, c) = i$. Continue.
- (3) If $(\exists e \in E | f_e = c \text{ and } l_e \geq q)$ then go to Step 4; otherwise, go to Step 7.
- (4) If $d = d_{\max}$ or $(\exists e \in E | f_e = c \text{ and } l_e = q)$ then go to Step 5; otherwise, go to Step 6.
- (5) Set $g = g + 1$, $v_g = c$, $w_g = \min\{l_e | e \in E, f_e = c, \text{ and } l_e \geq q\}$, $m(d, c) = g$, and $K_{n(d,c)m(d,c)} = 1$. Go to Step 7.
- (6) Set $m(d, c) = m(d + 1, c)$ and $k_{n(d,c)m(d,c)} = 1$. Continue.
- (7) Set $c = c + 1$, $q = c + d - 1$, $i = i + 1$, and $n(d, c) = i$. Continue.
- (8) If $(\exists e \in E | f_e \leq c \text{ and } l_e \geq q)$, $d = d_{\max}$, and $(\exists e \in E | f_e \leq c - 1 \text{ and } l_e = q - 1)$ then go to Step 10; otherwise go to Step 9.
- (9) If $(\exists e \in E | f_e = c \text{ and } l_e \geq q)$ and either $(\exists e \in E | f_e \leq c \text{ and } l_e = q)$ or $d = d_{\max}$, then go to Step 10; otherwise go to Step 11.
- (10) Set $g = g + 1$, $v_g = \max\{f_e | e \in E, f_e \leq c, \text{ and } l_e \geq q\}$, $w_g = \min\{l_e | e \in E, f_e \leq c, \text{ and } l_e \geq q\}$, $m(d, c) = g$, and $k_{n(d,c)m(d,c)} = 1$. Go to Step 14.
- (11) If $(\exists e \in E | f_e = c \text{ and } l_e \geq q)$ go to step 12; otherwise, if $(\exists e \in E | f_e \leq c \text{ and } l_e \geq q)$ go to Step 13; otherwise, go to Step 14.
- (12) Set $m(d, c) = m(d + 1, c)$ and $k_{n(d,c)m(d,c)} = 1$. Go to step 14.
- (13) If $(\exists e \in E | f_e \leq c - 1 \text{ and } l_e = q - 1)$ then set $m(d, c) = m(d + 1, c - 1)$; otherwise, set $m(d, c) = m(d, c - 1)$. In either case, set $k_{n(d,c)m(d,c)} = 1$ and go to Step 14.
- (14) If $c < P + 1 - d$, return to Step 7; otherwise, go to Step 15.
- (15) If $d > d_{\min}$, return to Step 2; otherwise, Stop.

APPENDIX B

A simple example is used to illustrate how the regions are utilized. For this example, assume that: (1) four employees are available, having the characteristics given in Table B1; (2) the operating day is nine hours; (3) planning intervals are hour-long; (4) shifts of four to eight hours are acceptable; and (5) no breaks are scheduled. For convenience, shift variables (t_{ij}) will be defined as the number of employees working a shift of i hours commencing in period j . Figure B1 illustrates a matrix of the shift variables and identifies the regions, while Table B2 provides additional information on the regions. For this example, constraint set (7) would be:

$$\begin{array}{rcl}
 y_{11} + y_{13} + y_{14} + y_{15} & \leq & 2 & \text{(employee 1)} \\
 & & y_{23} + y_{25} & \leq 1 & \text{(employee 2)} \\
 & & y_{32} + y_{34} + y_{35} & \leq 1 & \text{(employee 3)}
 \end{array}$$

Note that as employees 1 and 4 are available over identical periods, employee region variables are defined for only one of them (employee 1), and the right hand side of the set (7) constraint for that employee is set at 2. Constraint set (8) would be:

$$\begin{array}{rcl}
 t_{81} + t_{71} + t_{72} + t_{61} + t_{51} + t_{41} & = & y_{11} & \text{(region 1)} \\
 & & t_{73} + t_{64} + t_{55} + t_{46} & = y_{32} & \text{(region 2)} \\
 & & & & t_{62} + t_{52} + t_{42} & = y_{13} + y_{23} & \text{(region 3)} \\
 & & & & t_{63} + t_{54} + t_{45} & = y_{14} + y_{34} & \text{(region 4)} \\
 & & & & t_{53} + t_{43} + t_{44} & = y_{15} + y_{25} + y_{35} & \text{(region 5)} \\
 & & & & & & t_{82} & = 0 & \text{(region 6)}
 \end{array}$$

Note that any shifts which cannot be staffed (such as t_{82} in this example) may be eliminated from the model and the restriction defining the null region (region 6 in this example) could then be eliminated from constraint set (8).

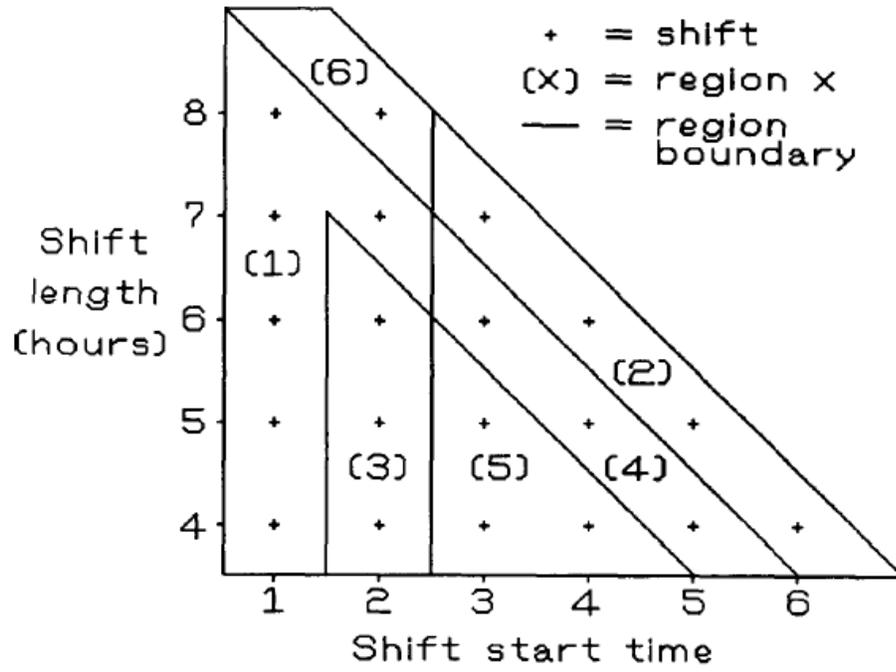
Table B1
THE AVAILABILITY OF FOUR EXAMPLE EMPLOYEES

Employee (e)	Earliest Possible Start Time (f_e)	Latest Possible Finish Time (l_e)
1	1	8
2	2	7
3	3	9
4	1	8

TABLE B2
INFORMATION ON THE REGIONS DEFINED BY THE
AVAILABILITY OF THE EXAMPLE EMPLOYEES

Region (g)	Earliest Shift Start (v_g)	Latest Shift Finish (w_g)	Employees Who Can Staff Shifts from the Region
1	1	8	1
2	3	9	3
3	2	7	1,2,4
4	3	8	1,3,4
5	3	7	1,2,3,4
6	1	9	none

FIGURE B1
THE SHIFT VARIABLES AND THE REGIONS
DEFINED BY THE AVAILABILITY OF THE EXAMPLE EMPLOYEES



REFERENCES

- Bailey, J. "Integrated Days Off and Shift Personnel Scheduling." *Computers and Industrial Engineering*, vol. 9, no. 4, 1985, 395-404.
- Baker, K.R., T.B. Crabill, and M.J. Magazine. "An Optimal Procedure for Allocating Manpower with Cyclic Requirements." *AIIE Transactions*, vol. 5, no. 2, June 1973, 119-126.
- Bechtold, S.E., and L.W. Jacobs. "Implicit Modeling of Flexible Break Assignments in Optimal Shift Scheduling." *Management Science*, vol. 36, no. 11, November 1990, 1339-1351.
- Dantzig, G.B. "A Comment on Edie's 'Traffic Delays at Toll Booths'." *Operations Research*, vol. 2, no. 3, 1954, 339-341.
- Glover, F, C. McMillan, and R. Glover. "A Heuristic Programming Approach to the Employee Scheduling Problem and Some Thoughts on 'Managerial Robots'." *Journal of Operations Management*, vol. 4, no. 2, 1984, 113-128.
- Henderson, W.B., and W.L. Berry. "Heuristic Methods for Telephone Operator Shift Scheduling: An Experimental Analysis." *Management Science*, vol. 22, no. 12, August 1976, 1372-1380.
- Keith, E.G. "Operator Scheduling." *AIIE Transactions*, vol. 11, no. 1, March 1979, 37-41.
- Loucks, J.S. "A Multiobjective Heuristic Procedure for the Tour-Scheduling/Task-Assignment of a Heterogeneous Workforce." In *Proceedings, 1986 National*

- Meeting of the Decision Sciences Institute*. Honolulu, HI: Decision Sciences Institute, November 1986.
- Love, R.R., Jr., and J.M. Hoey. "Management Science Improves Fast-Food Operations." *Interfaces*, vol. 20, no. 2, March-April 1990, 21-29.
- Moondra, S.L. "An L.P Model for Work Force Scheduling for Banks." *Journal of Bank Research*, vol. 7, no. 4, 1976, 299-301.
- Morris, J.G., and M.J. Showalter. "Simple Approaches to Shift, Days-Off and Tour Scheduling Problems." *Management Science*, vol. 29, no. 8, August 1983, 942-50.
- Segal, M. "The Operator-Scheduling Problem: A Network-Flow Approach." *Operations Research*, vol. 22, no. 4, 1974, 808-823.
- Thompson, G.M. "A Multi-Phase Heuristic for Developing Shift Schedules Using Non-Continuously Available, Homogeneously-Skilled Employees. ' ' Working paper 90-1, College of Business, University of Utah, Salt Lake City, UT, 1990.