

Labor Staffing and Scheduling Models for Controlling Service Levels

Gary M. Thompson

Cornell University

June 16, 1997

Forthcoming in *Naval Research Logistics*

Correspondence address

Gary M. Thompson
Operations Management
School of Hotel Administration
Cornell University
Ithaca, NY 14853

Acknowledgement

The author wishes to acknowledge the helpful comments of two anonymous referees and an Associate Editor.

Abstract

The problems of labor staffing and scheduling have received substantial attention in the literature. We introduce two new models of the labor staffing and scheduling problems that avoid the limitations of existing models. Collectively, the models have five important attributes. First, both models ensure the delivery of a minimally acceptable level of service in all periods. Second, one model can identify the least expensive way of delivering a specified aggregate level of customer service (the labor staffing problem and a form of labor scheduling problem). Third, the other model can identify the highest level of service attainable with a fixed amount of labor (the other form of the labor scheduling problem). Fourth, the models enable managers to identify the pareto relationship between labor costs and customer service. Fifth, the models allow a degree of control over service levels that is unattainable with existing models. Because of these attributes, which existing models largely do not possess, we expect these models to have broad applicability in a wide range of organizations operating in both competitive and noncompetitive environments.

Labor Staffing and Scheduling Models for Controlling Service Levels

INTRODUCTION

Labor staffing attempts to identify a workforce's ideal size and composition. Its planning horizon is medium-duration and a function, among other things, of the amount of time required to hire and train new employees. Labor scheduling has a much shorter planning horizon, often only a week or even a day. Labor scheduling, which considers the workforce size and composition as givens, has two common forms. In the higher-flexibility form, employees are not guaranteed a set number of hours over the planning horizon while, in the lower-flexibility form, employees are guaranteed a set number of hours over the planning horizon. These types of situations are usually seen with workforces largely comprised of part-time and full-time employees, respectively.

Shift scheduling—developing a single day's work schedule—and tour scheduling—developing a weekly work schedule—are common manifestations of the labor staffing and scheduling problems. Our focus in this paper is primarily upon situations where the interests of the customer and the server are not mutual, making it difficult or impossible to assign a monetary value to the cost of waiting [42]. In such environments, management often wishes (a) to deliver a specified level of service as inexpensively as possible (the problems of labor staffing and scheduling a flexible workforce) or (b) to maximize the level of service delivered by a fixed labor cost or pool (the problem of scheduling an inflexible workforce). As did others [1, 14, 23, 28, 40, 45], we measure *service* as the percentage of customers served within a set number of minutes after their arrival.

In this paper we make a distinction between the aggregate threshold service level—the overall level of service that management wishes to provide customers—and the minimum

acceptable service level—the lowest level of customer service that management considers acceptable in any planning period. For example, management may wish to serve at least 80% of all customers arriving during a week within 2 min of their arrival—the threshold level of service—yet want every planning period to serve at least 50% of customers within 3 min of their arrival—the minimum acceptable service level.

The distinction between aggregate threshold and minimum acceptable service levels is new to the labor scheduling literature. It originates from our observations of how managers apply existing labor scheduling models. In traditional labor scheduling models, each planning period's staffing requirement is set at the minimum number of staff who can deliver what we choose to call the threshold level of service. We have often observed that managers are willing to reduce the staffing levels in some periods if they can measurably reduce the overall labor scheduled. We have observed that this reduction is typically performed without any particular regard to the effects on service that the staff reduction might have. The advantage of our new models is that they deliver a specified minimum level of service in each period, and ensure the delivery of an overall level of service. For example, using the terminology of [22], the minimum acceptable service level could be set at the service “qualifier” level, while the overall service level could be set at the service “winner” level. Our models are thus applicable to a wide range of organizations, including those operating in competitive environments.

Two simple examples should serve to clarify the nature of the environments upon which we focus. First, regulated monopolies, such as local telephone service providers, commonly have externally imposed regulations requiring they deliver specified aggregate threshold levels of service in certain customer service functions [33]. Second, consider a customer service department of a consumer products firm and the issues involved in staffing its toll-free,

telephone call center. Since the customers are typically calling for information about the product or to voice a complaint regarding the product—actions with which it is difficult to associate a direct monetary effect of good or poor service—the firm will experience difficulty in determining the ideal staffing levels from an economic perspective. However, the firm can use focus groups and simple experiments to identify the level of service that customers deem appropriate for this function. Further, they can solicit information from customers in two forms: What is the typical level of service the customers would like to experience (which, for our purposes, is the aggregate threshold service level) and what (lower) level of service is minimally acceptable to the customers (for our purposes, the minimum acceptable service level) . These service levels are inputs of the models we develop.

Although an extensive literature exists on tour and shift labor staffing and scheduling, most researchers have used formulations based on the models of [15] or [27]. As we will discuss later, these models have limitations in problem environments like the ones we consider. In this paper, we develop two new models of the labor staffing and scheduling problems. The new models, by using more complete information on the merits of different staff sizes and by simultaneously determining both the number of employees needed in each period and the labor schedule, avoid the limitations of existing models. They also allow a much higher degree of managerial control over service levels than is possible with existing models. The new models also allow one to identify the pareto frontier between service cost and service level. Knowing where one's competitors are on the service cost/service level frontier—and something about customer choice patterns—should enable managers to make better choices about where to position their service on the frontier.

In the remainder of the paper, we provide useful background information on labor scheduling and identify limitations of existing labor staffing and scheduling models, introduce the new labor staffing and scheduling models, describe two simulation experiments designed to evaluate the usefulness of the new models, and present and discuss the results of these experiments. The paper concludes by offering suggestions for future research.

LABOR SCHEDULING BACKGROUND

In this section we present definitions of terms, sets, constants, and variables used in the classic labor staffing and scheduling models, followed by a summary of these models. Lastly, we identify limitations both of these models and of the labor scheduling literature, in general.

Definitions of Terms, Indices, Sets, Constants, and Variables

Define the following terms:

Period—a basic interval for which planning occurs, typically 15, 30, or 60 min.

Shift—a work schedule for an employee for a specific day, which observes applicable restrictions defined by management policy, contractual obligations, or governmental regulations.

Tour—a work schedule for an employee for a week, comprised of separate shifts, which satisfies applicable restrictions defined by management policy, contractual obligations, or governmental regulations.

Service level—the percentage of customers served within a specified waiting time limit.

Aggregate threshold service level—the percentage of customers arriving over the duration of the planning horizon (a week for the tour scheduling problem, a day for the

shift scheduling problem) that management wishes to have served within the specified waiting time limit.

Desired staffing level (for a period)—the minimum number of staff that can deliver the aggregate threshold level of service within the period.

Minimum acceptable service level—the lowest level of service that is acceptable in any planning period.

Minimum acceptable staffing level (for a period)—the minimum number of staff that can deliver the minimum acceptable level of service within the period.

Define the following sets:

P = the set of planning periods,

T = the set of allowable tours (shifts).

Define the following constants:

$$a_{tp} = \begin{cases} 1, & \text{if period } p \text{ is a working period for tour (shift) } t, \\ 0, & \text{otherwise,} \end{cases}$$

b_p^β = the limit on the bounded shortage of employees in period p ,

b_p^π = the limit on the bounded surplus of employees in period p ,

c_t = the cost of having an employee work tour (shift) t ,

c^α = the cost of the unbounded shortage of employees, per employee-period of shortage,

c^β = the cost of the bounded shortage of employees, per employee-period of shortage,

c^π = the cost of the bounded surplus of employees, per employee-period of surplus,

c^θ = the cost of the unbounded surplus of employees, per employee-period of surplus,

r_p = the desired staffing level in period p .

Define the following variables:

α_p = the unbounded shortage of employees in period p ,

β_p = the bounded shortage of employees in period p ,

π_p = the bounded surplus of employees in period p ,

θ_p = the unbounded surplus of employees in period p ,

x_t = the number of employees working tour (shift) t .

Notable Tour and Shift Labor Staffing and Scheduling Models

Dantzig [15] presented the original mathematical programming formulation of the tour and shift labor scheduling problems (DLSM) :

$$\text{Min } Z = \sum_{t \in T} c_t x_t \quad (1)$$

subject to

$$\sum_{t \in T} a_{tp} x_t \geq r_p \quad \text{for } p \in P, \quad (2)$$

$$x_t \geq 0 \text{ and integer for } t \in T. \quad (3)$$

DLSM's objective (1) is to minimize the cost of the scheduled tours (shifts), subject to the restriction, constraint set (2), that sufficient employees are present in all periods.

Keith [27] developed a model (KLSM) having two employee shortage and surplus variables per planning period, which consequently allows fewer than the desired number of staff to be scheduled in any period:

$$\text{Min } Z = \sum_{t \in T} c_t x_t + \sum_{p \in P} (c^\alpha \alpha_p + c^\beta \beta_p + c^\pi \pi_p + c^\theta \theta_p), \quad (4)$$

subject to

$$\sum_{t \in T} a_{tp} x_t + \alpha_p + \beta_p - \pi_p - \theta_p = r_p \quad \text{for } p \in P, \quad (5)$$

$$\beta_p \leq b_p^\beta \quad \text{for } p \in P, \quad (6)$$

$$\pi_p \leq b_p^\pi \quad \text{for } p \in P, \quad (7)$$

$$\alpha_p, \beta_p, \pi_p, \theta_p \geq 0 \text{ and integer for } p \in P, \quad (8)$$

and (3)

KLSM's objective includes both the costs of tours (shifts) and the cost of deviations from the desired staffing levels. Constraint set (5) measures, for each period, deviations from the desired staffing levels using two shortage and two surplus staffing variables. One shortage and one surplus staffing variable from each period have upper bounds; these variables also cost less than the corresponding unbounded variables. Thus, solutions to KLSM typically have actual employee shortages and surpluses falling within the specified bounds.

Limitations of Labor Staffing and Scheduling Models

Begin by considering how to set the desired staffing levels in DLSP to meet management's objective of delivering an aggregate threshold level of customer service as inexpensively as possible. First, assume the staff size is such that every period delivers at least the aggregate threshold level of service. In this case, the actual aggregate level of service will be greater than preferred, even if the staff levels are satisfied exactly in every period. The reason for this is that since staffing levels must be increased by one employee at a time, the smallest number of employees who deliver at least the aggregate threshold level of service almost always deliver better than the aggregate threshold level of service. Next, assume that each period's employee requirement is set so as to deliver the minimum acceptable level of service. In this case, nothing ensures the delivery of the aggregate threshold service level.

If, in the former case, a schedule has substantial overstaffing, we have observed that managers commonly relax the employee requirements in the binding periods, those having no

overstaffing, and develop a new schedule based on the adjusted staffing requirements. By doing this, managers are, in essence, implementing a model more similar to KLSM. Relaxing DLSP's staffing levels and redeveloping a schedule has two outcomes. First, there is nothing to ensure that the resultant schedule provides the aggregate threshold level of service. Second, even if the resultant schedule does provide the aggregate threshold level of service, there is no way of determining short of a very exhaustive trial-and-error search if the schedule offers the least expensive way of doing so. In the latter case, if the service level falls below the aggregate threshold level, a manager may decide to increase the employee requirements in some periods and redevelop a schedule. Even if the resultant schedule then delivers the aggregate threshold service level, one cannot readily tell if it is the least expensive way of doing so.

KLSM suffers from similar difficulties, since even if each period's desired staff size delivers the aggregate threshold level of service, nothing ensures that the aggregate service level equals or exceeds the threshold level. One might think that KLSM's objective function, which includes costs of employee shortages and surpluses, would, by balancing the periods providing better-than-threshold service against the periods providing poorer-than-threshold service, deliver the aggregate threshold service level. For this balancing to happen correctly, however, a model needs a measure of the true impact of staff size changes on customer service, not a surrogate measure of that impact. LIM1 shall refer to the dilemma of setting staffing levels that deliver the aggregate threshold service level as inexpensively as possible.

A limitation specific to KLSM, LIM2, is its failure to ensure the delivery of the minimum acceptable service level in every period. This failing occurs because there is no limit on the amount of understaffing that can occur in any period when using KLSM. One might argue that the cost of the unbounded understaffing variables could be set prohibitively high. If the costs are

high enough, the unbounded understaffing variables could be removed from KLSM since they would always be zero, anyway, but this changes the nature of the model. Now consider how to maximize the level of service provided by a fixed workforce. A limitation of DLSM in this case, LIM3, is its assumption that surplus employees are of equal value, regardless of the period in which the surplus falls. Although KLSM uses two-tiered costs for under- and overstaffing, its costs for each tier are equal within and across periods, and so LIM3 applies also to KLSM. LIM3 becomes apparent when one examines organizations' service delivery functions, and sees that these functions are typically queuing systems. A well-known characteristic of queuing systems is the nonlinear, diminishing impact on system performance of adding service personnel. LIM3 means that a model considers all feasible schedules of equal cost to be equally good, even when, in reality, they are not. Thus, maximizing the level of service with a fixed labor cost or labor pool is problematic for a model subject to LIM3.

To summarize the limitations of existing models in the environments we consider, then: LIM1 is the inability of existing models DLSM and KLSM to deliver an aggregate threshold service level as inexpensively as possible. LIM2 is KLSM's failure to ensure that some minimal number of staff, delivering a minimally acceptable level of service, is present each period. Finally, LIM3 means the neither DLSM or KLSM can maximize any measure of customer service with a fixed workforce. The implication of these limitations is that the existing models do not guarantee an efficient deployment of labor. Consequently, a firm using existing models will be unable to position itself on the pareto frontier of labor costs and threshold service levels.

Table 1 categorizes pertinent tour and shift labor staffing and scheduling research. Although most researchers have used models based on DLSM, some researchers have used models similar to KLSM but without the bounded employee shortage and surplus variables. With

two exceptions [48, 49], the entire labor scheduling literature is susceptible to LIM3. LIM2 applies for all models based on KLSM, with one exception [4]. With three exceptions, the entire literature is susceptible to LIM1. The exceptions [29, 34, 36] all examine the scheduling of bank staff for the processing of checks, a quasi-manufacturing environment. Since it is not necessary to process all checks arriving in a period, this environment is unusual in that it can carry an inventory of unprocessed checks. Unfortunately, for services where the item being processed is customers, carrying an inventory of unprocessed customers is often not a realistic option. The near universal occurrence of the limitations we have noted—and the difficulty these limitations pose for managers—motivated our development of improved tour and shift labor staffing and scheduling models. The next section introduces these models.

TOUR/SHIFT MODELS ALLOWING CONTROL OVER SERVICE LEVELS (CLSM)

We shall refer to the new models generically as CLSM for Controlled Labor Scheduling Model. CLSM-CST, the first model, is used for delivering a managerially specified aggregate threshold service level as inexpensively as possible. Because CLSM-CST contains no restrictions on the amount of labor available or that must be scheduled, it applies both for the labor staffing problem and for the labor scheduling problem in those cases where employees are not guaranteed a minimum amount of work over the scheduling horizon. CLSM-CST is

Min 1

subject to

$$\sum_{t \in T} a_{tp} x_t - \sum_{j=1}^{w_p} \mu_{pj} \geq m_p \text{ for } p \in P, \quad (9)$$

$$\sum_{p \in P} \sum_{j=1}^{w_p} n_{pj} \mu_{pj} \geq \Gamma - \sum_{p \in P} \delta_p, \quad (10)$$

$$\mu_{pj} = \{0,1\} \text{ for } p \in P \text{ and } j = 1, \dots, w_p \quad (11)$$

and (3),

m_p = the minimum acceptable staff size in period p ,

$$\mu_{pj} = \begin{cases} 1, & \text{if the number of employees working in period } p \text{ equals or exceeds } m_p + j, \\ 0, & \text{otherwise} \end{cases}$$

w_p = the number of employees in excess of the minimum acceptable staff size for period p who, when added to the schedule, incrementally increase the expected number of customers served within the designated limit on waiting time (that is, $n_{p,w_p} > 0$),

n_{pj} = the incremental increase in the expected number of customers served within the designated waiting time limit, occurring with the addition of the $(m_p + j)$ th employee in period p ,

δ_p = the expected number of customers served within the designated waiting time limit, with the minimum acceptable number of employees working in period p ,

Γ = the expected number of customers who must be served within the designated limit on waiting time over the duration of the operating week (day) to meet the aggregate threshold level of service.

As with DLSM, CLSM-CST's objective includes only tour (shift) costs. CLSM-CST measures, for each period, the surplus of employees over the minimum acceptable staffing level in constraint set (9). Since the smallest possible staff size in each period is the minimum acceptable number of employees for the period (i.e., the number of staff who can deliver the minimally acceptable level of service), constraint set (9) enables CLSM-CST to avoid LIM2. Constraint (10) ensures the delivery of the threshold aggregate level of service *over the duration of the operating week (day)*, thus enabling CLSM-CST to avoid LIM1. LIM3 does not apply to CLSM-CST because the binary variables measuring employee surpluses (the μ_{pj}) each have an

associated customer-service coefficient. These coefficients allow CLSM-CST to represent the nonlinear, declining effect on customer service of adding personnel in each period.

The second new model, CLSM-SRV, addresses the labor scheduling problem where the workforce size or labor cost is fixed. Here, management wishes to maximize the overall level of service subject to a limitation on the amount or cost of labor while still delivering the minimum acceptable level of service in each planning period. CLSM-SRV is

$$\text{Max } \sum_{p \in P} \sum_{j=1}^{w_p} n_{pj} \mu_{pj}, \quad (12)$$

subject to

(3), (9), (11), and

$$\sum_{t \in T} c_t x_t \leq \text{cost}_{max} \quad (13a)$$

and/or

$$\sum_{t \in T} x_t \leq \text{labor}_{max} \quad (13b)$$

where cost_{max} is the maximum amount management wishes to spend on labor and labor_{max} is the maximum number of tours (shifts) to be scheduled. Note that constraint (13a) will be selected if management wishes to limit the schedule's cost, while constraint set (13b) will be selected if management wishes to restrict the number of employees scheduled. Also note that the left sides of equations (13a) and (13b) will tend to approach the maximum value (the right sides), given the objective of maximizing customer service. LIM1 does not apply in the type of problem CLSM-SRV considers, while LIM2 and LIM3 do not apply for CLSM-SRV for the same reasons they do not apply to CLSM-CST. An unabridged version of this paper [50] presents examples of DLSM, KLSM, CLSM-CST, and CLSMSRV for a simple scheduling scenario. In the experiments described in the next section we employ two forms of CLSM-SRV. CLSM-SRV(D)

and CLSM-SRV(K) maximize the aggregate threshold service while respectively limiting the labor cost to that found by DLSM and KLSM.

DESIGN OF THE MODEL VALIDATION EXPERIMENTS

CLSM is more complex in form than either DLSM or KLSM. We developed two model validation experiments, both simulating service delivery systems, to evaluate whether CLSM's complexity is an asset or a liability and to gauge the benefits that one might expect when implementing CLSM in practice. In the following subsections, we identify experimental assumptions, describe the experimental designs, address model-related issues, and discuss information accuracy.

Experimental Assumptions

Experiment 1 addressed shift scheduling for an operating day comprised of 36 30-min planning periods. Experiment 2 addressed tour scheduling for seven operating days, each comprised of 18 hour-long planning periods. Other experimental assumptions are: the true, or underlying daily and weekly customer arrival-rate patterns are stationary, but unknown; the mean daily or weekly customer arrival rate is constant; both inter-arrival times and service times follow exponential distributions; there is a single, first-come-first-served queue with, typically, multiple service personnel; no customers balk from the service system; and a minimum acceptable level of service requiring the presence of enough staff to ensure every period's customer service rate exceeded its expected customer arrival rate.

Experimental Designs

Each experiment had a full-factorial design. Because of the exploratory nature of the experiments, we selected factors, and levels of these factors, potentially affecting the relative performance of the labor scheduling models. As Table 2 shows, independent variables in the experiments were the daily customer arrival-rate pattern, the weekly, across-day variation in the mean daily customer arrival rate (Experiment 2 only); the variation of the customer arrival-rate pattern, the mean duration of customer service, the aggregate threshold level of customer service, the level of flexibility existing when developing the labor schedule (Experiment 1 only), the variability in shift or tour costs, and, finally, the model used in developing the labor schedule. The combinations of arrival patterns and arrival rate variation yielded eight distinct daily customer arrival-rate curves in Experiment 1, as Figure 1 shows. In Experiment 2, the combinations of customer arrival patterns, relative daily arrival rates, and arrival pattern variability yielded 16 distinct weekly customer arrival-rate curves. Additional details on the experimental factors and the simulation process may be found in the unabridged paper [50].

We expect CLSM to perform better than DLSM and KLSM across all levels of the experimental factors. Thus, our null hypotheses are that CLSM-SRV(D) and CLSM-CST have the same performance as DLSM and that CLSM-SRV(K) and CSLM-CST have the same performance as KLSM. In other words, our first null hypothesis is that CLSM-SRV(D) and CLSM-CST both deliver the same level of service at the same cost as DLSM. Our second null hypothesis is that CLSM-SRV(K) delivers the same level of service at the same cost as KLSM. The alternate hypotheses are that CLSM performs better; i.e., it delivers the aggregate threshold service level less expensively or provides a higher level of service at the same or lower labor cost.

Model-Related Issues

The final experimental factor, the labor scheduling model, had five levels: DLSM, KLSM, and CLSM (three implementations). Below we describe the employee requirements used in the models, describe our implementation of KLSM and the three versions of CLSM, and describe the heuristic process used for generating solutions to the models.

Employee Requirements in the Models. As is traditional, the employee requirement for each period in DLSM and KLSM was determined from each period having to deliver the threshold service level. Thus the desired staff size in each period for these models was the smallest number of staff that could serve 50% (Experiment 1) or 75% (Experiment 2) of the expected customers within the specified waiting time limit, a number determined with an $M/M/s$ queuing model. To account for the spillover of customer service across planning periods, we used a procedure [46] to adjust each period's recorded arrival rate.

Each period's minimum acceptable staff size was the same in all versions of CLSM, and equal to the smallest number of staff whose total service rate exceeded the expected arrival rate of customers in the period. An $M/M/s$ queuing model provided the customer-service coefficients associated with CLSM's binary overstaffing variables, after adjusting each period's recorded arrival rate using a procedure to address the effect of service spillover across planning intervals [46].

KLSM Implementation Issues. In our implementation of KLSM, the per-employee-period costs of the bounded (c^π) and unbounded (c^θ) overstaffing variables and of the bounded (c^β) and unbounded (c^α) understaffing variables were 0.0, 0.5, 2.0, and 8.0 employee periods, respectively, in Experiment 1 and 0.0, 0.5, 3.0, and 10.0 employee-period, respectively, in

Experiment 2. Although arbitrary, these values are consistent with the suggestions of Keith [27].

The following formula set the limits on the bounded under- and overstaffing variables in KLSM:

$$b_p^\beta = b_p^\pi = \min\{4, \max\{1, \lfloor .2r_p + .5 \rfloor\}\}, \quad (14)$$

where

$$\lfloor x \rfloor = \text{the smallest integer } \geq x.$$

As noted earlier, since KLSM does not restrict the maximum shortage of employees that can occur in any period, it fails to ensure the delivery of the minimum acceptable level of service in all periods. To actually ensure this requires that one modify KLSM to limit the maximum amount of understaffing that can occur in each planning period. We thus implemented a modified version of KLSM that restricts the understaffing variables in each planning period using Eq. (14) and

$$\alpha_p + \beta_p \leq r_p - m_p \quad \text{for } p \in P \quad (15)$$

Three Versions of CLSM. We investigated three implementations of CLSM, one based on CLSM-CST and two based on CLSM-SRV (CLSM-SRV(D) and CLSM-SRV(K)). To allow us to gauge CLSM's value in problems where the objective is delivering an aggregate threshold level of service as inexpensively as possible, CLSM-CST had an aggregate threshold service level requiring the service of 50% (Experiment 1) or 75% (Experiment 2) of the customers within the designated waiting time limit, the same threshold service level used in setting the staffing levels in DLSP and KSLM.

The limit on labor in CLSM-SRV(D) and CLSM-SRV(K), imposed using Eq. (13a), was the actual cost of labor scheduled in the solutions to DLSP and KLSM, respectively. Thus, comparing the service levels provided by CLSM-SRV(D) (CLSM-SRV(K)) to those provided by

DLSM (KLSM) gives one a measure of the benefit of using CLSM for the problem of labor scheduling with a fixed labor cost (or labor pool).

Solving the Models Heuristically. We used a heuristic solution procedure when developing each model's labor schedules, because the sheer number of test problems made optimally generating schedules impractical. Our heuristic was based on the best heuristic for labor scheduling problems yet developed, the simulated annealing heuristic of [13]. As we discuss in more detail later, we ran the annealing heuristic for 60 s for each model on each problem. We coded both labor scheduling heuristics and the service delivery simulation in FORTRAN. Experiment 1 heuristics and simulation were compiled using Microsoft FORTRAN Powerstation and run on a Pentium-90 based personal computer under MSDOS 6.2. Experiment 2 heuristics and simulation were compiled using Microsoft FORTRAN Powerstation for Windows NT and run on a Pentium-100 based personal computer in a DOS window using the Microsoft Windows NT operating system.

To establish the effectiveness of our annealing heuristic, we randomly selected four problems from Experiment 2 and compared the annealing solutions to solutions to DLSM, KLSM, and CLSM-CST obtained using the commercial integer programming (IP) software combination of GAMS [12] and OSL [24]. We ran each problem for 10 min using the IP software, and recorded the best solution found during that time if the problem had not solved to verified optimality. We call the solutions found using the IP software the "best known" solutions. Over the four problems, our annealing heuristic for DLSM (KLSM) found two (two) optimal solutions, had an average solution cost 0.11% (0.10%) above the best known solutions, and had a worst case performance of 0.28% (0.33%) more costly than the best known solution. In all four problems, our annealing heuristic for CLSM-CST found better solutions in 1 min than

the best solution found in 10 min using the IP software. Overall, the performance of our annealing heuristics compare favorably to earlier implementations of simulated annealing (annealing yielded the best solutions on 94% of test problems [13]; average solutions found using simulated annealing were 0.29% more costly than optimal [49]). We should note that we performed no specialized tuning of GAMS/OSL, and so it is possible that these results may overestimate the heuristic's performance. Our previous experience with GAMS/OSL [48, 49] has been favorable, however, reducing this concern.

Information Accuracy

The labor scheduling process involves developing a labor schedule to satisfy expected customer demand as cheaply as possible. Expected customer demand in each planning period is an input to an appropriate simulation or queuing model (M/M/s for our experiment) for determining the desired and minimum acceptable staffing levels and for determining the customer-service coefficients associated with CLSM's binary overstaffing variables. If CLSM only outperformed DLSSM and KLSM in an ideal world, where actual customer arrivals equaled expected arrivals, CLSM would be of limited value. The assumptions of stationary underlying customer arrival-rate curves (as illustrated in Fig. 1) and exponentially distributed inter-arrival and service times enabled us to generate expected customer arrivals by simulating historical arrivals to the service delivery system. Simulating five days or weeks of past customer arrivals yielded the expected numbers of customers arrivals, by period, and the expected customer service durations. Obviously, there will be greater (smaller) variation between each period's observed and true arrival rates when the observed mean is determined using less (more) historical information.

RESULTS OF THE MODEL VALIDATION EXPERIMENTS

Table 3 summarizes the experimental results relating to labor costs and actual service levels. Overall on Experiment 1 (Experiment 2), CLSM-CST yielded cost savings of 9.40% (18.97%) compared to DLSP's schedules and 2.76% (0.52%) compared to KLSM's schedules. Importantly, CLSM's lower average schedule cost did not come at the expense of a lower-than-wished-for level of customer service: On average, CLSM's schedules in Experiment 1 (Experiment 2) served 14.04% (8.20%) *more* customers within the specified waiting time limit than required by the aggregate threshold service level.

Identifying a pareto frontier between labor costs and aggregate threshold service levels can help managers position their services relative to those of their competitors. To investigate the pareto frontier, we randomly selected three problems from Experiment 2 and used CLSM-SRV to examine the aggregate threshold level of service possible under increasing labor costs. Figure 2 illustrates the frontiers for these problems and, for comparison, the solutions to DLSP and KLSM. Although the illustrated frontiers were based on solutions obtained using the simulated annealing heuristic, our testing of the heuristic suggests that the frontiers, as presented, closely represent the true frontiers. Figure 2 shows that DLSP's and KLSM's solutions are in the interior of the frontiers, and hence not efficient alternatives. Examining KLSM's solutions shows two of the outcomes possible with this model. Figures 2(a) and 2(b) show KLSM's solutions to be respectively more and less expensive than CLSM-CST's solutions (solutions may also be equally expensive). One might ask, why can KLSM have a solution less costly than CLSM-CST, since CLSM-CST is supposed to deliver the aggregate threshold service level as inexpensively as possible? The reason, as careful examination of Figure 2(b) shows, is that KLSM failed to deliver the aggregate threshold service level in this problem. This can happen with KLSM

because there is nothing in the model to ensure that the desired level of service is provided (as per our discussion of LIM1).

We observed four interesting results when we examined model performance across the levels of the experimental factors. First, CLSM-CST's schedules had lower average costs than those of DLSM and KLSM, for all levels of all experimental factors in both experiments. Second, CLSM-CST's schedules delivered better-than-threshold customer service, on average, for all levels of all experimental factors in both experiments. Third, the cost advantage of CLSM-CST's schedules increased markedly in both experiments as the limit on customers' waiting time decreased from 2 to 1.8 min. Finally, only with the Unimodal arrival-rate pattern in Experiment 1 did CLSM-SRV fail to deliver a higher level of service at a lower cost than DLSM. Here, CLSM-SRV served 2.65% more customers within the specified waiting time while using only 0.04% more employee-periods of labor (398.452 employee periods for CLSM-SRV, on average, versus 398.288 employee periods for DLSM).

DISCUSSION

In this section we present a discussion addressing the selection of the best model, why CLSM works, and the use of CLSM in other environments.

Choosing the Best Overall Model

It is useful to examine employee utilization before making a final judgment on the superiority of one model over another. For example, if a less costly model increased employee utilization dramatically, which in turn led to overworked employees and higher turnover, a model may be less attractive than it would appear only from a labor cost perspective. As shown

in Table 3, DLSP, KLSM, and CLSM-CST had actual employee utilizations of approximately 70.5%, 75.7%, and 77.8%, respectively, in Experiment 1 and utilizations of approximately 48.4%, 59.4%, and 59.7%, respectively, in Experiment 2.

We judge KLSM as dominating DLSP, since it delivered the desired aggregate threshold service level less expensively, and deployed labor in a much more efficient way [see, from Fig. 2, how much closer KLSM is to the pareto frontier than is DLSP and, from Table 3, how much closer KLSM is to CLSM-SRV(K) than DLSP is to CLSM-SRV(D)]. Nonetheless, we judge CLSM as dominating KLSM for three reasons. First, CLSM-CST delivers the desired aggregate threshold service level less expensively than does KLSM. Second, CLSM-SRV deploys its labor more effectively than does KLSM (KLSM is close to, but not on the pareto frontier defined by CLSM-SRV). Third, KLSM offers far less control over the service levels than does CLSM.

Both forms of CLSM performed their intended functions. CLSM-CST, designed to find the cheapest way of delivering an aggregate threshold level of service, delivered the aggregate threshold level of service in Experiment 1 (Experiment 2) at a cost of 2.76% and 9.40% (18.97% and 0.52%) lower than KLSM and DLSP, respectively. Only a moderate increase in employee utilization occurred with CLSM-CST, mitigating any concerns about increased turnover resulting from overworked employees. CLSM-SRV(D) in Experiment 1 (Experiment 2) served 2.61% (0.86%) more customers within the desired time at a cost approximately 0.1% less (0.05% less) than DLSP. The only exception to CLSM-SRV(D)'s complete dominance over DLSP was at the Unimodal arrival pattern in Experiment 1, where CLSM-SRV(D) delivered a notably higher level of service at a slightly higher cost. This anomaly is likely a minor artifact of the simulation process, in which the actual labor usage may be slightly higher than that scheduled in advance.

Thus, the results offer compelling evidence of the correctness of our initial hypotheses of CLSM's superiority.

A key advantage of CLSM is that one can use it for either the labor staffing or the labor scheduling problem. CLSM-CST is applicable when labor is variable. This is the case in the labor staffing problem—which seeks to find the best size of the workforce—and in the labor scheduling problem where employees are not guaranteed a set number of hours each week (or day). CLSM-SRV, on the other hand, is most valuable for scheduling a fixed workforce who are guaranteed a set number of hours each week (day). Another benefit of CLSM is that it one can use it to identify the pareto relationship between service and labor requirements for specified minimum acceptable and aggregate threshold service levels, as illustrated in Figure 2. The pareto frontier clearly shows the inferiority of KLSM and DLSM, but more important, the pareto frontier should help managers identify the level of service to provide customers, and the cheapest way of doing so. As stated earlier, knowing where one's competitors are on the service cost/service level frontier — and something about customer choice patterns—should enable managers to make better choices about where to position one's service on the frontier.

A final advantage of CLSM-CST is that its cost advantage over DLSM increased as the aggregate threshold level of customer service increased. In Experiment 1 (Experiment 2), decreasing the limit on customers' waiting time from 2 to 18 min increased CLSM-CST's labor cost advantage over DLSM from 7.90% to 10.83% (16.68% to 23.02%) and from 2.11% to 3.40% (0.07% to 0.96%) over KLSM. To understand why CLSM-CST's advantage increases at higher service levels, one must recognize that the minimum acceptable staff size in a period can deliver a low level of customer service. Obviously then, if the minimum acceptable and desired staff sizes are identical in each period—which will occur with low levels of threshold customer

service—then DLSM and CLSM-CST will yield schedules of identical cost (though CLSM-SRV can be used to maximize the level of service delivered at this cost) . Having CLSM-CST's relative advantage over DLSM and KLSM increase with higher levels of threshold customer service is an attractive characteristic given that good service is important in the long-run success of service organizations [16].

Why CLSM Works

CLSM performed better than DLSM and KLSM for two reasons: First, CLSM determines the number of staff needed in each planning period simultaneously with the development of the labor schedule, and second, CLSM incorporates information on the non-linear value of adding staff in each period. That CLSM performed as well relative to KLSM and DLSM with relatively little historical information suggests that CLSM does not depend on inherently superior information for its superior performance, but rather its superior performance comes from an inherently superior way of using whatever information is available.

Using CLSM in Other Environments

Though we selected a traditional M/M/s queuing environment to validate the superior utility of CLSM (an environment having Poisson arrivals, exponential service times, no balking, a single waiting line), one should not draw the conclusion that this is the only environment in which CLSM has superior utility. To apply CLSM in other service environments only requires that one calculate the binary overstaffing variables' customer-service coefficients in a manner appropriate to the particular environment. For example, CLSM can also readily be applied in environments relaxing any or all of the assumptions in our test environment.

CONCLUSIONS

This paper presented two new models of the labor staffing and scheduling problems (CLSM-CST and CLSM-SRV). Advantages of the new models include: 1) CLSM-CST delivers an aggregate threshold level of service more cheaply than DLSM ([15]'s 1954 model) or KLSM (a modified form of [27]'s 1979 model), which is important for labor staffing and for scheduling a workforce that can work variable hours; 2) CLSM-SRV delivers better service than DLSM or KLSM with the same amount of labor, which is important for scheduling a workforce guaranteed hours; 3) both deliver the minimum acceptable level of service in all planning periods, while the original form of KLSM [27] does not; 4) they do what astute managers attempt to do with DLSM (and what KLSM was designed to do)—balance periods of better than desired service against periods with lower than desired but still acceptable service—but they do it much easier (one need only solve CLSM once); 5) they provide managers an easy means of identifying the pareto relationship between customer service and labor costs for their service delivery system, which should help firms deliver the correct level of service at the best price.

The results of our experiment suggest several avenues for additional research. First, a reevaluation of scheduling procedures reported in the literature would be useful. Many reported techniques would require modification for use with the nonlinear value of employees inherent in CLSM. Without modification, these techniques could only develop schedules for a less-desirable formulation of the labor scheduling problem. Thus, to improve the practical value of labor scheduling research, new scheduling methods should be evaluated and existing scheduling methods should be reevaluated using one or both forms of CLSM. Fortunately, heuristic solution procedures such as simulated annealing, which have yielded excellent performance on labor scheduling problems [13, 49], are readily adaptable to solve CLSM.

Second, an interesting study would evaluate the effect on labor costs and customer service of varying the minimum acceptable level of customer across various threshold levels of customer service. (In effect, this would be the development of a three-dimensional pareto frontier. Figure 2 shows only a two-dimensional pareto frontier, since the minimum acceptable level of customer service was constant in our experiment.) Obviously, if the minimum acceptable level of customer service is identical with the threshold service level, CLSMCST would have no cost advantage over DLSP, since each period's minimum acceptable staff size in CLSM-CST would equal the desired staffing levels in DLSP. Even in this environment, however, CLSM-SRV can be used to maximize the level of service using a specified amount of labor.

REFERENCES

- [1] Andrews, B.H., and Parsons, H.L., “L.L. Bean Chooses a Telephone Agent Scheduling System,” *Interfaces*, 19(6), 1–9 (1989).
- [2] Bailey, J., “Integrated Days Off and Shift Personnel Scheduling,” *Computers and Industrial Engineering*, 9(4), 395–404 (1985).
- [3] Bailey, J., and Field, J., “Personnel Scheduling with Flexshift Models,” *Journal of Operations Management*, 5(3), 327–338 (1985).
- [4] Baker, K.R., “Workforce Allocation in Cyclical Scheduling Problems: A Survey,” *Operational Research Quarterly*, 27(1), 155–167 (1976).
- [5] Baker, K.R., Crabill, T.B., and Magazine, M.J., “An Optimal Procedure for Allocating Manpower with Cyclic Requirements,” *AIIE Transactions*, 5(2), 119–126 (1973).
- [6] Bartholdi, J.J., “A Guaranteed-Accuracy Round-Off Algorithm for Cyclic Scheduling and Set Covering,” *Operations Research*, 29(3), 501–510 (1981).
- [7] Bechtold, S.E., Brusco, M.J., and Showalter, M.J., “A Comparative Evaluation of Labor Tour Scheduling Methods,” *Decision Sciences*, 22(4), 683–699 (1991).
- [8] Bechtold, S.E., and Jacobs, L.W., “Implicit Optimal Modeling of Flexible Break Assignments in Labor Staffing Decisions for Service Operations,” *Management Science*, 36(11), 1339–1351 (1990).
- [9] Bechtold, S.E., and Jacobs, L.W., “Improvement of Labor Utilization in Shift Scheduling for Services with Implicit Optimal Modeling,” *International Journal of Operations and Production Management*, 11(2), 54–69 (1991).
- [10] Bechtold, S.E., and Showalter, M.J., “Simple Manpower Scheduling Methods for Managers,” *Production and Inventory Management*, 26(3), 116–132 (1985).

- [11] Bechtold, S.E., and Showalter, M.J., “A Methodology for Labor Scheduling in a Service Operating System,” *Decision Sciences*, 18(1), 89–107 (1987).
- [12] Brooke, A., Kendrick, D., and Meeraus, A., GAMS, Release 2.25, A User’s Guide, *The Scientific Press*, South San Francisco, CA, 1992.
- [13] Brusco, M.J., and Jacobs, L.W., “A Simulated Annealing Approach to the Cyclic Staff-Scheduling Problem,” *Naval Research Logistics*, 40, 69–84 (1993).
- [14] Buffa, E.S., Cosgrove, M.J., and Luce, B.J., “An Integrated Work Shift Scheduling System,” *Decision Sciences*, 7(4), 620–630 (1976).
- [15] Dantzig, G.B., “A Comment on Edie’s ‘Traffic Delays at Toll Booths,’ ” *Operations Research*, 2(3), 339–341 (1954).
- [16] Davis, M., “How Long Should a Customer Wait for Service?” *Decision Sciences*, 22(2), 421–434 (1991).
- [17] Easton, F.F., and Rossin, D.F., “Sufficient Working Subsets for the Tour Scheduling Problem,” *Management Science*, 37(11), 1441–1451 (1991a).
- [18] Easton, F.F., and Rossin, D.F., “Equivalent Alternate Solutions for the Tour Scheduling Problem,” *Decision Sciences*, 22(5), 985–1007 (1991b).
- [19] Gaballa, A., and Pearce, W., “Telephone Sales Manpower Planning at Qantas,” *Interfaces*, 9(3), 1–9 (1979).
- [20] Henderson, W.B., and Berry, W.L., “Heuristic Methods for Telephone Operator Shift Scheduling: An Experimental Analysis,” *Management Science*, 22(12), 1372–1380 (1976).
- [21] Henderson, W.B., and Berry, W.L., “Determining Optimal Shift Schedules for Telephone Traffic Exchange Operators,” *Decision Sciences*, 8(2), 239–255 (1977).
- [22] Hill, T., *Manufacturing Strategy*, Irwin, Homewood, IL, 1989.

- [23] Holloran, T.J., and Byrn, J.E., “United Airlines Station Manpower Planning System,” *Interfaces*, 16(1), 39–50 (1986).
- [24] IBM Corporation, Optimization Subroutine Library, Release 2, Kingston, NY, 1991.
- [25] Jacobs, L.W., and Bechtold, S.E., “Labor Utilization Effects of Labor Scheduling Flexibility Alternatives in a Tour Scheduling Environment,” *Decision Sciences*, 24(1), 148–166 (1993).
- [26] Jarrah, A.I.Z., Bard, J.F., DeSilva, A.H., “Solving Large-Scale Tour Scheduling Problems,” *Management Science*, 40(9), 1124–1144 (1994).
- [27] Keith, E.G., “Operator Scheduling,” *AIIE Transactions*, 11(1), 37–41 (1979).
- [28] Kolesar, P.J., Rider, K.L., Crabill, T.B., and Walker, W.E., “A Queuing-Linear Programming Approach to Scheduling Police Patrol Cars,” *Operations Research*, 23(6), 1045–1062 (1975).
- [29] Krajewski, L.J., Ritzman, L.P., and McKensie, P., “Shift Scheduling in Banking Operations: A Case Application,” *Interfaces*, 10(2), 1–8 (1980).
- [30] Li, C., Robinson, E.P., Jr., and Mabert, V.A., “An Evaluation of Tour Scheduling Heuristics with Differences in Employee Productivity and Cost,” *Decision Sciences*, 22(4), 700–718 (1991).
- [31] Loucks, J.S., and Jacobs, F.R., “Tour Scheduling and Task Assignment of a Heterogeneous Work Force: A Heuristic Approach,” *Decision Sciences*, 22(4), 719–738 (1991).
- [32] Love, R.R., Jr., and Hoey, J.M., “Management Science Improves Fast Food Operations,” *Interfaces*, 20(2), 21–29 (1990).
- [33] Lynch, J.G., Jr., Buzas, T.E., and Berg, S.V., “Regulatory Measurement and Evaluation of Telephone Service Quality,” *Management Science*, 40(2), 169–194 (1994).

- [34] Mabert, V.A., “A Case Study of Encoder Shift Scheduling under Uncertainty,” *Management Science*, 25(7), 623–631 (1979).
- [35] Mabert, V.A., and Showalter, M.J., “Measuring the Impact of Part-Time Workers in Service Organization,” *Journal of Operations Management*, 9(2), 209–229 (1990).
- [36] Mabert, V.A., and Watts, C.A., “A Simulation Analysis of Tour-Shift Construction Procedures,” *Management Science*, 28(5), 520–532 (1982).
- [37] McGinnis, L.F., Culver, W.D., and Deane, R.H., “One- and Two-Phase Heuristics for Workforce Scheduling,” *Computers and Industrial Engineering*, 2(1), 7–15 (1978).
- [38] Moondra, S.L., “An L.P. Model for Work Force Scheduling for Banks,” *Journal of Bank Research*, 7(4), 299–301 (1976).
- [39] Morris, J.G., and Showalter, M.J., “Simple Approaches to Shift, Days-Off and Tour Scheduling Problems,” *Management Science*, 29(8), 942–950 (1983).
- [40] Segal, M., “The Operator-Scheduling Problem: A Network-Flow Approach,” *Operations Research*, 22(4), 808–823 (1974).
- [41] Showalter, M.J., and Mabert, V.A., “An Evaluation of a Full-/Part-Time Tour Scheduling Methodology,” *International Journal of Operations and Production Management*, 8(7), 54–71 (1989).
- [42] Taha, H.A., “Queuing Theory in Practice,” *Interfaces*, 11(1), 43–49 (1981).
- [43] Taylor, P.E., and Huxley, S.J., “A Break from Tradition for the San Francisco Police: Patrol Officer Scheduling Using an Optimization-Based Decision Support System,” *Interfaces*, 19(1), 4–24 (1989).
- [44] Thompson, G.M., “Shift Scheduling When Employees Have Limited Availability: An L.P. Approach,” *Journal of Operations Management*, 9(3), 352–370 (1990).

- [45] Thompson, G.M., “Improving the Utilization of Front-Line Service Delivery System Personnel,” *Decision Sciences*, 23(5), 1072–1098 (1992).
- [46] Thompson, G.M., “Accounting for the Multi-Period Impact of Service When Determining Employee Requirements for Labor Scheduling,” *Journal of Operations Management*, 11(3), 269–287 (1993).
- [47] Thompson, G.M., “Improved Implicit Optimal Modeling of the Labor Shift Scheduling Problem,” *Management Science*, 41(4), 595–607 (1995).
- [48] Thompson, G.M., “Labor Scheduling Using NPV Estimates of the Marginal Benefit of Additional Labor Capacity,” *Journal of Operations Management*, 13(1), 67–86 (1995).
- [49] Thompson, G.M., “A Simulated-Annealing Heuristic for Shift Scheduling Using Noncontinuously Available Employees,” *Computers and Operations Research*, 23(3), 275–288 (1996).
- [50] Thompson, G.M., “Labor Staffing and Scheduling Models for Controlling Service Levels: Unabridged Version,” Working paper, School of Hotel Administration, Cornell University, 1996.
- [51] Vohra, R.V., “A Quick Heuristic for Some Cyclic Staffing Problems with Breaks,” *Journal of the Operational Research Society*, 39(11), 1057–1061 (1988).

Table 1. A classification of pertinent labor staffing and scheduling literature.

Reference	Scheduling model ^a	Applicable limitations ^b		
		LIM1	LIM2	LIM3
[3], [5–11], [13], [15], [17–21], [25], [26], [30], [31], [35], [38–41], [45], [47], [51]	DLSM	Yes	No	Yes
[27], [29] ^c , [34] ^c , [36] ^c	KLSM	Yes	Yes	Yes
[2], [14], [32], [37] ^d , [43], [44]	KLSM-NBV	Yes	Yes	Yes
[48], [49]		Yes	No ^e	No

^a The model closest in form to that used in the reference. DLSM = the model of [15]. KLSM = the model of [27]. NBV = a model allowing unbounded employee shortages and surpluses but lacking variables representing bounded employee shortages and surpluses.

^b LIM1 = the problem of setting the employee requirements to meet the aggregate threshold service level as inexpensively as possible. LIM2 = failing to ensure that each period delivers the minimum acceptable level of customer service. LIM3 = an inability to maximize the level of service provided by a fixed workforce, because the model assumes that a surplus employee is of equal value, regardless of the period in which the surplus falls.

^c These researchers examined the quasimanufacturing problem of scheduling bank employees to process checks. LIM1 does not apply to these works because inventory (unprocessed checks) may be carried in this environment.

^d [37] apparently allowed understaffing, but measured performance only using idle time (overstaffing).

^e The models developed by [48] and [49] would, with a minor adjustment, avoid LIM2.

Table 2. Overview of the independent variables in the service delivery system simulation experiments.

Factor (number of levels)	Levels	Used in experiment
Customer arrival-rate pattern (4)	Unimodal (1 daily peak), bimodal (2 daily peaks), trimodal (3 daily peaks), and random (many daily peaks)—see Figure 1	Both
Variation in the customer arrival-rate pattern (2)	Coefficients of variation in the customer arrival-rate pattern of 0.25 and 0.50 (sine function amplitudes of 0.353 and 0.706)	Both
Variation in the mean daily customer arrival rates	Mean daily arrival rates of 1.0, 0.9, 0.9, 1.0, 1.0, 1.1, and 1.1 and 1.0, 0.8, 0.8, 1.0, 1.0, 1.2, and 1.2 for Sunday through Saturday, respectively	Exp. 2
Mean service duration (2)	2 and 16 min per customer	Both
Aggregate threshold level of customer service (2)	Serving 50% (Experiment 1) or 75% (Experiment 2) of customers within 1/8 and 2 min of their arrival	Both
Labor scheduling flexibility (2)	95 and 515 different shifts available when developing the labor schedule	Exp. 1
Shift/tour cost variability (2)	Shifts (Experiment 1) or tours (Experiment 2) with costs, per working period, randomly assigned from the ranges of 1 { 0.0, and 1 { 0.1	Both
Labor scheduling model (5)	DLSM, KLSM, and 3 versions of CLSM—CLSM-CST, CLSM-SRV(D), and CLSM-SRV(K)	Both


Table 3. Labor cost and service level results.^a

Experiment	Model	Actual labor costs ^b	Actual service level ^c	Actual labor utilization
One (addressing daily shift scheduling)	CLSM-SRV(D)	485.98	81.27	70.59
	DLSM	459.45	79.20	70.51
	CLSM-SRV(K)	427.93	71.89	75.71
	KLSM	428.11	71.62	75.68
	CLSM-CST	416.28	64.04	77.83
Two (addressing weekly tour scheduling)	CLSM-SRV(D)	2341.48	96.35	48.43
	DLSM	2342.54	95.53	48.40
	CLSM-SRV(K)	1907.49	85.17	59.45
	KLSM	1908.05	85.01	59.43
	CLSM-CST	1898.17	83.20	59.74


^a All values are averages across the 384 simulated service environments.

^b Measured in employee-period-equivalents.


^c Average proportion of customers served within the designated maximum waiting time. The aggregate threshold level of service required serving 50% of the customers within the appropriate waiting time limit in Experiment 1 and 75% of customers in Experiment 2.



c. Low-variation bimodal customer arrival pattern.



e. Low-variation trimodal customer arrival pattern.



g. Low-variation random customer arrival pattern.







Figure 1. Daily customer arrival-rate curves in the experiments: (a) low-variation unimodal customer arrival pattern. (*continues*)



d. High-variation bimodal customer arrival pattern.



f. High-variation trimodal customer arrival pattern.



h. High-variation random customer arrival pattern.




Figure 1 (continued). (b) High-variation unimodal customer arrival pattern.




Figure 2. Pareto frontiers for schedule cost and aggregate threshold service levels for two randomly selected problems in Experiment 2. (Note that all tours are equally costly in these two problems.)

(a) First randomly selected problem. (continues)




Figure 2 (continued) . (b) Results from the second randomly selected problem from Experiment 2.