IIE Transactions (2005) **37**, 175–187

ISSN: 0740-817X print / 1545-8830 online
DOI: 10.1080/07408170590885657

Equipment scheduling at mail processing and distribution centers

XINHUI ZHANG[1*] and JONATHAN F. BARD[2]

[1]*Department of Biomedical, Industrial and Human Factor Engineering, Wright State University, Dayton, OH 45435, USA*
E-mail: xinhui.zhang@wright.edu
[2]*Graduate Program in Operations Research & Industrial Engineering, The University of Texas, ETC 5.160, C2200, Austin, TX 78712-1063, USA*
E-mail: jbard@mail.utexas.edu

Received November 2002 and accepted April 2004

Mail Processing and Distribution Centers (P&DCs) run 24 hours a day, 7 days a week and are staffed by a skilled complement of full-time, part-time, and temporary employees. A recurrent problem faced by facility managers involves the development of daily schedules for the automation equipment used to process the mail. The ultimate goal is to minimize the labor costs associated with running the facility while ensuring that all service standards are met. The focus of this paper is on the equipment scheduling aspect of the problem. In particular, we seek a weekly schedule that satisfies all operational, technological, and legal constraints of the system at a minimum cost. The problem is modeled as a large-scale mixed-integer linear program and solved sequentially using a three-stage methodology. In each stage, a separate criterion is optimized and the corresponding objective function value is used as a constraint in subsequent stages. To ease the computational burden, two major enhancements are developed. The first is a pre-processor designed to reduce the number of integer variables; the second is a heuristic that uses the linear programming solution as a target and attempts to find a feasible integer point as close to it as possible. The methodology is demonstrated with data obtained from the Dallas P&DC. The computations indicate that for letter operations alone the annual savings will be on the order of $1.6 million per facility when the system is implemented nationwide over the next 3 years.

1. Introduction

The United States Postal Service (USPS) operates a network of approximately 275 major mail Processing and Distribution Centers (P&DCs) that serve as the interfaces between local post offices and the rest of the nation. The amount of mail passing through these facilities is staggering. On a typical day, a medium-sized P&DC might receive as many as 5000 000 letters, 500 000 flats, and thousands of parcels. To handle such a large volume, advanced equipment in the form of optical character readers, automated facer-cancellers, barcode sorters, and material handling systems have been especially designed to automate as much of the mail stream as possible. Today, approximately 95% of letters and flats are sequenced automatically in the order in which they will be delivered by the carriers.

P&DCs are like high volume factories and run 24 hours a day, 7 days a week. They are staffed by a skilled workforce comprising full-time, part-time, and casual employees whose weekly tours are subject to labor laws, union contracts, and local policies. The problem of scheduling the automation equipment in a P&DC with as few workers as possible is a challenge that the USPS has yet to meet. For a given equipment schedule, however, the problem of determining the optimal size of the workforce has been solved, and is currently being implemented nationwide (Jarrah *et al.*, 1994a; Bard *et al.*, 2003). The purpose of this paper is to describe a modeling approach that has been developed to provide optimal equipment schedules for use with the staff scheduler.

Ideally, we would like to solve both problems at once but this has not been possible due to the overwhelming complexity of the operational environment. Nevertheless, our experience has shown that for both equipment and labor scheduling, a hierarchical approach works best. At the first stage, one or more large-scale mixed-integer programs is solved to obtain facility-wide schedules. In subsequent stages, post-processors are used to transform these schedules into daily and hourly staff and machine assignments.

The primary contribution of this paper centers on the methodology developed to solve the mixed flow shop–job shop problem that arises in high volume factories. For such problems, there are at least two complications that have been largely overlooked by the research community:

*Corresponding author

(i) there is no single objective that can be used to guide the analysis; and (ii) machines must be operated by skilled personnel. The latter is more critical to mail processing facilities. The ultimate objective is to minimize the cost of running a facility, which, for a given equipment configuration, reduces to minimizing the cost of the workforce when the capital costs are sunk. Nevertheless, without proper management of both staffing and equipment considerations, suboptimal solutions may result.

This issue is taken up further in the next section. In Section 3, we give an overview of mail processing at a P&DC, followed in Section 4 by our sequential approach to the equipment scheduling problem. The idea is to solve a series of mixed-integer linear programs, each designed with a different aspect of the operational environment in mind. The solution framework is discussed in Section 5. In Section 6, we present our computational experience with the proposed methodology using data from the Dallas P&DC. To highlight the results, in Section 7 we compare the staffing costs imposed by our schedule with the costs imposed by the schedule in use by the facility at the time of the study. Finally, concluding remarks are given in Section 8.

2. Background and scheduling issues

The equipment scheduling problem takes as input mail volume and arrival profiles, and determines the "optimal" use of equipment over the day (Bard *et al.*, 1993; Jarrah *et al.*, 1994b; Berman *et al.*, 1997). The resultant schedule and assignment of operations to machines is the front-end of the staff scheduling problem. Given an equipment schedule in terms of demand for workers, the staff scheduling problem is: (i) to find the minimum size workforce needed to run the facility; and (ii) to construct weekly tours for all regular employees that comply with contractual rules and policies (Malhotra *et al.*, 1992; Beaumont, 1997; Brusco and Jacobs, 1998). The two components of the joint problem and their relationships are shown in Fig. 1.

Determining even a rough relationship between an equipment schedule and the corresponding workforce requirements is no easy task, if for no other reason than the difference in the planning horizon for the two problems. Equipment schedules are generated daily to match the incoming mail whereas staff schedules are generated weekly due to union rules and the nature of the job. Rather than trying to solve the problems jointly, a more manageable approach is to investigate the characteristics of different equipment schedules, determine the staffing requirements they impose, and then select a schedule that minimizes the cost of the accompanying workforce.

Along these lines, we have identified several criteria for directing the search for high quality equipment schedules. The foremost of these is an approximation of the daily staffing cost. Others include the need to process as much of the arriving mail as possible on the day it arrives, to minimize machine startups, and to compress the schedule as much as possible. These criteria are further discussed in Section 4 where the full model is presented.

The first efforts to analyze mail processing facilities were aimed at selecting an optimal mix of equipment rather than generating optimal schedules. Early studies by Bard *et al.* (1993) and Jarrah *et al.* (1994b) were undertaken to support the acquisition of automation equipment. In the most comprehensive, Jarrah *et al.* (1994b) presented a Mixed-Integer Linear Program (MILP) to solve the equipment selection problem in a P&DC. An auxiliary linear program was used as a post-processor to derive tentative flows through the facility to ensure efficient machine utilization.

Berman *et al.* (1997) used mail processing centers as an example of high volume factories and modelled it as a network of workstations. They focused mainly on the costs associated with different full-time to part-time labor ratios and used linear programming for the analysis.

Jarrah *et al.* (1994a) developed an integrated methodology for solving the P&DC staff scheduling problem. Their model combined the daily shift scheduling problem with the days-off assignment problem in a novel way that allowed for weekly tour construction. User options included variable shift starting times, full-time to part-time labor ratios, and lunch breaks. Bard *et al.* (2003) provided a number of enhancements that paved the way for implementation. The methodology presented in this paper, is similarly aimed at elevating the sophistication of the previous work to the point where implementation is possible.

3. Overview of a P&DC

Mail arriving at a P&DC can be categorized as follows: letters, flats (e.g., large envelopes and magazines), bundles, parcels, and packages. Only letters are considered in this paper. Flats are governed by similar processing concepts and

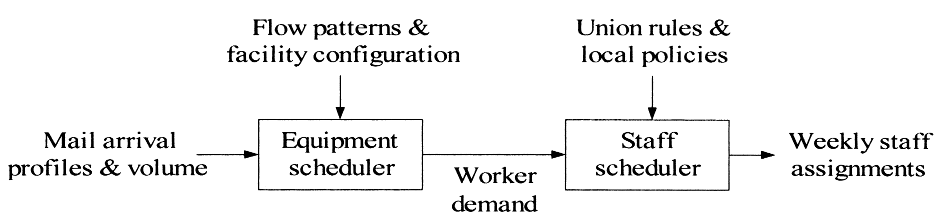

Fig. 1. The USPS scheduling problem.

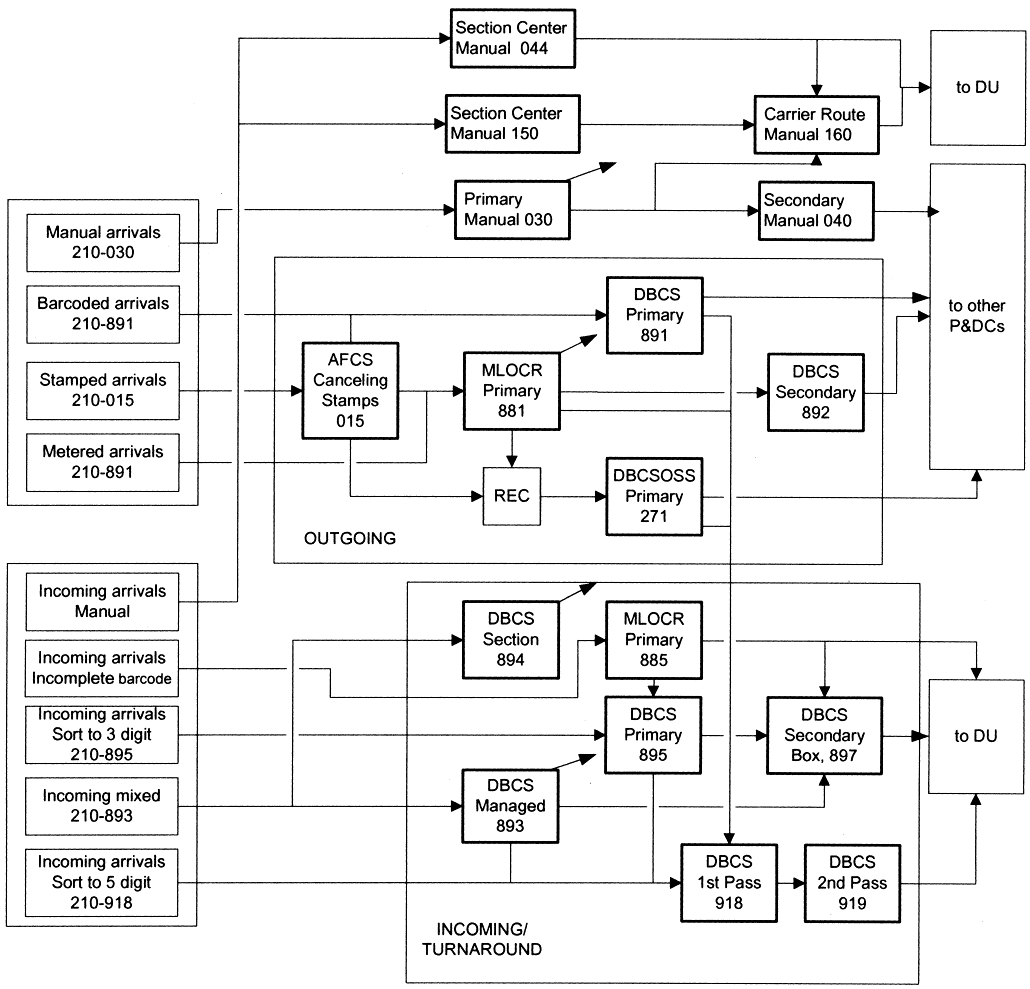

Fig. 2. Major mail flows in a P&DC.

are handled the same way, but with different equipment. The other categories are either not suited for automation or too small in volume to be of concern.

The automated processing of a letter follows three steps: (i) canceling the stamp (if one exists); (ii) reading the address and identifying the destination with a barcode; and (iii) sorting the letter to its final destination. Four types of equipment are currently deployed in P&DCs for letter processing: (i) an advanced face-canceller system (AFCS); (ii) a multi-line optical character reader (MLOCR); (iii) a remote barcode sorter (RBCS); and (iv) a delivery barcode sorter (DBCS). The functionality of the RBCS is often provided by a delivery barcode sorter fitted with an output subsystem (DBCS-OSS) in these facilities. Figure 2 displays the major operations and mail flows associated with letters at the Dallas facility. Arrows represent mail flows and blocks represent the processing operations (the identification numbers and the equipment used for these operations are also included). An operation is considered major if it must be run for at least 10 hours a day. A mail flow is considered major if the output to the destination node consists of more than 15% of the total output at the originating node. Similar diagrams exist for flats, parcels, and bundles.

4. Problem formulation

4.1. *Basic multi-level lot sizing model*

Processing activities at a P&DC can be modeled as a multi-level lot sizing problem with setup times for each day of the week. Here, multi-level refers to the fact that the input of an operation depends not only on external arrivals, but also on flows from upstream operations. In general, machines are not set up and run unless an equivalent of at least 30 minutes of a particular mail type is available for processing. For this reason, we divide the day into 48 30-minute periods.

In the model, it is assumed that: (i) incoming mail arrives at the beginning of each period and is available for immediate processing; (ii) mail processed in one time period will not be available for downstream operations until the next time period to account for transportation time between machines; (iii) a certain amount of time is required to set up and clear a machine between operations; and (iv) only one operation can be performed on a machine in a single period. These assumptions are consistent with current practice.

With this background in mind, we now present the basic multi-level lot sizing model for the equipment scheduling problem. The following notation is used in the developments.

Indices

i, o = indices for input and output mail streams;
n, p = indices for nodes (a unique operation is associated with each node);
g = index for machine groups;
t = index for time periods; $t \in T = \{1, \ldots, 48\}$, $t = 1$ corresponds to 7:00 a.m.

Sets

G = set of machine groups, $G = \{$AFCS, MLOCR, DBCS, DBCS-OSS$\}$;
$G(n)$ = set of machine groups capable of performing the operation at node n;
I, O = sets of input and output mail streams for the facility;
N = set of nodes;
$N(g)$ = set of nodes whose operations can be performed by machines in group g;
$P(n)$ = set of nodes immediately preceding node n;
$I(n)$ = set of input mail streams to node n;
$O(n)$ = set of output mail streams from node n;
$T(n)$ = set of periods during which an operation at node n can be performed;
$T(i)$ = set of periods during which mail arrivals i may be processed; generally, $T(i) \subseteq T(n)$ where n is the node that accepts the arrivals of input stream i.

Parameters

u_n = amount of mail left unprocessed at node n from the previous day;
$a_i(t)$ = amount of external mail of type $i \in I$ arriving at the facility during period t, $t \in T(i)$;
$m_g(t)$ = number of machines available in group g during period t;
ρ_n = processing rate for operation at node n (pieces/period);
f_{npn} = fraction of mail processed at predecessor node p that is sent to successor node n;
τ_1 = the amount of time required to start up a machine; currently set at 10 minutes;
τ_2 = the amount of time required to clear a machine; currently set at 10 minutes.

Decision variables

$v_n(t)$ = inventory level of mail at node n at the end of period t, $t \in T$;
$w_{ng}(t)$ = amount of mail processed at node n by machine group g during period t, $t \in T(n)$;
$Y_{ng}(t)$ = number of machines devoted to the operation that is performed at node n by machine group g during period t, $t \in T(n)$;
$Z^1_{ng}(t)$ = number of startups in machine group g for operation n at the beginning of period t, $t \in T(n)$;
$Z^2_{ng}(t)$ = number of clearance operations performed at node n in machine group g at the end of period t, $t \in T(n)$.

Constraints

$$u_n + \sum_{i \in I(n), 1 \in T(i)} a_i(1) - \sum_{g \in G(n)} w_{ng}(1) = v_n(1), \quad \forall n \in N, \tag{1a}$$

$$v_n(t-1) + \sum_{i \in I(n), t \in T(i)} a_i(t) + \sum_{p \in P(n)} f_{pn} \sum_{g \in G(p)} w_{pg}(t-1) - \sum_{g \in G(n)} w_{ng}(t) = v_n(t), \quad \forall t \in T \backslash \{1\},\ n \in N, \tag{1b}$$

$$w_{ng}(t) + \frac{\tau_1}{30}\rho_n Z^1_{ng}(t) + \frac{\tau_2}{30}\rho_n Z^2_{ng}(t) \le \rho_n Y_{ng}(t), \quad \forall g \in G(n),\ t \in T(n),\ n \in N, \tag{1c}$$

$$\sum_{n \in N(g)} Y_{ng}(t) \le m_g(t), \ \forall t \in T,\ g \in G, \tag{1d}$$

$$Z^1_{ng}(t) - Z^2_{ng}(t-1) = Y_{ng}(t) - Y_{ng}(t-1), \quad \forall t \in T(n),\ n \in N(g),\ g \in G, \tag{1e}$$

$$Z^1_{ng}(t) \ge 0,\ Z^2_{ng}(t) \ge 0,\ \forall t \in T(n),\ n \in N(g), g \in G, \tag{1f}$$

$$v_n(t), w_{ng}(t) \ge 0,\ \forall t \in T(n),\ n \in N,\ g \in G(n), \tag{1g}$$

$$Y_{ng}(t) \ge 0 \text{ and integer}, \quad \forall g \in G(n), t \in T(n),\ n \in N. \tag{1h}$$

1. *Inventory balance constraints of Equations* (1a) *and* (1b). Constraint (1a) stipulates that for each operation n at the beginning of the day in period 1, the mail remaining from the previous day, plus the sum of external arrivals, minus the sum of mail processed during this period, must equal the ending inventory indicated on the right-hand side. Constraint (1b) imposes the same requirements in all other periods but with the inclusion of the mail transferred from predecessor nodes.
2. *Production capacity limitation of Equation* (1c). Constraint (1c) stipulates that in period t, the workload at node n allocated to group g, plus the lost capacity during startup and clearance, must be less than or equal to the processing capacity of the machines in group g assigned to operation n at time t.
3. *Machine capacity limitation of Equation* (1d). Constraint (1d) ensures that the number of machines in group g operating in time period t is less than or equal to the number of machines available.

4. *Definition of startup and clearance activities in Equation* (1e). Constraint (1e) defines the startup and clearance activities, each of which takes a certain amount of time denoted by τ_1 and τ_2, respectively. As mentioned, the lost production capacity associated with these activities is taken into account in constraint (1c).
5. *Constraints on variables in Equations* (1f), (1g) *and* (1h). Constraints (1f), (1g) and (1h) define the non-negative values of all variables and the integer requirement of variable $Y_{ng}(t)$. Notice that the startup and clearance variables are not defined as being integer. As we will see, because we are trying to minimize a linear function of the startup variables, for a given integer value of $Y_{ng}(t)$, $Z^1_{ng}(t)$ and $Z^2_{ng}(t)$ will be integer-valued in any feasible solution.
6. *Additional flow management considerations*. To provide more control over the pattern of work flow over the day, it might be desirable to add a constraint originally proposed by Berman *et al.* (1997). In specifying the constraint, let $\theta_n(t')$ be the fraction of work associated with operation $n \in N' \subseteq N$ that must be completed by time $t' \in T(n)$. The following constraint ensures that the volume of mail processed before time t' (left-hand side) is greater than or equal to $\theta_n(t')$ of the mail processed over the day:

$$\sum_{t \le t'} \sum_{g \in G(n)} w_{ng}(t) \ge \theta_n(t') \sum_{t \in T(n)} \sum_{g \in G(n)} w_{ng}(t), \quad \forall t' \in T(n),\ n \in N'. \tag{1i}$$

The inclusion of Equation (1i) allows the user to better manage the workflow of each operation over the course of the day. However, it requires a somewhat arbitrary specification of the fractions $\theta_n(t')$ for each $t' \in T(n)$ and $n \in N'$. Poor choices could seriously over-constrain the model and lead to inferior, if not infeasible, results. In light of this, our recommendation is to omit Equation (1i) in the first run and, if the solutions are not satisfactory, add it in subsequent runs to achieve a more desirable workflow distribution.

4.2. *Optimality criteria*

Four criteria have been identified for evaluating an equipment schedule: (i) the percentage of arriving mail processed in a day; (ii) the number of 8-hour shifts required to run the equipment; (iii) the number of machine setups in a day; and (iv) the weighted sum of the number of working periods in a day. The second criterion represents an estimate of staffing costs and is the most important. The first criterion represents a commitment to service self-imposed by the USPS, the third and fourth criteria offer further opportunities to improve the equipment schedules. These criteria are explained below.

4.1.1. *Minimize the ending inventory: process as much mail as possible*

The fundamental requirement of an equipment schedule is to process all the mail arriving at a facility in a timely manner. Due to capacity limitations of the equipment, however, this may not be possible so some of the mail may have to be held over to the next day. To push as much mail as the constraints permit, we solve the following problem for each day of the week.

$$\text{Minimize} \sum_{n \in N} v_n(48) \text{ subject to Equations (1a)–(1h)} \qquad \text{(objective 1).}$$

The objective is to minimize the ending inventory. Conservation of flow at each node in the network guarantees that there will be no accumulation of mail at any upstream workstation unless the capacity of the system limits its processing.

4.1.2. *Minimize the number of full-time shifts: an approximation to the daily staffing cost*

As previously mentioned, the predominant cost in running a P&DC is associated with the workforce. Maintenance is minimal by comparison, and the equipment is viewed as a sunk cost. The quality of an equipment schedule, then, must be judged by the staffing cost it imposes. As such, we introduce the second criterion: the number of full-time shifts. This number represents an approximation of the daily staffing cost and, as shown presently, provides a link between equipment scheduling and staff scheduling.

A shift is defined as a set of continuous time periods within a day, and is characterized by a start time and a length. At P&DCs, a full-time shift spans $8^1/_2$ hours including a half-hour lunch break. As a simplification, we only consider full-time shifts and treat them as continuous working periods of 8 hours with predetermined start times, but without lunch breaks. Because different skills are required to operate the equipment, it is necessary to associated shifts with machine groups. The following additional notation is used to model this aspect of the problem.

Indices

k = index for worker categories;
f = index for shifts.

Sets

K = set of worker categories;
F = set of shifts;
$G(k)$ = set of machine groups a worker in category k can operate.

Parameters

s_f = starting period for shift f;
e_f = ending period for shift f;
r_{kg} = number of workers of type k required to run a machine in group g.

Variables

ω_{kf} = number of workers of type k assigned to shift f.

Partial model

$$\text{Minimize} \quad \sum_{k\in K}\sum_{f\in F}\omega_{kf}, \qquad \text{(objective 2)}$$

subject to

$$\sum_{f:s_f\le t\le e_f}\omega_{kf} \ge \sum_{g\in G(k)}\sum_{n\in N(g)} r_{kg}Y_{ng}(t), \quad \forall t\in T,\ k\in K, \qquad \text{(2a)}$$

$$\omega_{kf}\ge 0 \text{ and integer}, \quad \forall f\in F, \quad k\in K. \qquad \text{(2b)}$$

The objective is to minimize the number of shifts. Constraint (2a) ensures that the number of workers on duty in period t (left-hand side) is sufficient to match the number requested (right-hand side), whereas constraint (2b) restricts ω_{kf} to be integer valued. The full problem includes the addition of constraints (1a)–(1h).

4.1.3. *Minimize the number of startups and clearances*

The third criterion used to further judge an equipment schedule is the number of startups and clearance activities. The objective is:

$$\text{Minimize} \quad \sum_{n\in N}\sum_{g\in G(n)}\sum_{t\in T(n)}\left(Z^1_{ng}(t)+Z^2_{ng}(t)\right), \quad \text{(objective 3a)}$$

where $Z^1_{ng}(t)$ and $Z^2_{ng}(t)$ are the number of startup and clearance operations. Minimizing the sum of startups and clearance activities reduces the number of times that the machines are switched on and off as well as the number of transitions from one operation to another. The result should lead to a more leveled and continuous schedule. Note that the number of startups equals the number of clearances so it is not necessary to include both terms in the above objective function.

4.1.4. *Minimize the weighted sum of the number of working periods per operation*

The fourth criterion is the weighted sum: the number of machines periods used to process the mail. Before we present this model, let us first looked at the unweighted sum: the total number of working periods in the day:

$$\text{Minimize} \quad \sum_{n\in N}\sum_{g\in G(n)}\sum_{t\in T(n)} Y_{ng}(t).$$

Given the volume of mail associated with an operation, the smaller the value of this objective function, the more likely machines will be running close to their capacity.

In practice, it is also preferable to shorten the working interval of an operation (the duration of an operation from start to end). A slight modification of the above objective function can help achieve this goal. In particular, we add a coefficient to $Y_{ng}(t)$ that decreases with time. The new problem is to minimize the weighted sum of working periods:

$$\text{Minimize} \quad \sum_{n\in N}\sum_{g\in G(n)}\sum_{t\in T(n)}(1-\varepsilon t)Y_{ng}(t), \quad \text{(objective 3b)}$$

where $\varepsilon > 0$ is a "small" number lying, say, between (0, 0.02], and $t \in [1, 48]$. When t is small, the coefficient $(1-\varepsilon t)$ is relatively large (close to one) so processing in earlier periods is penalized more than processing in later periods. This has the effect of pushing an operation to as late in the operation window as possible, as well as implicitly shortening the working intervals. The closer the value of ε is to 0.02, the stronger the effect. Testing has shown that a value of $\varepsilon = 0.01$, mid-way in the range (0, 0.02] provides good results.

Setting ε to a value above 0.02 (actually $1/48 = 0.0208$) would result in some coefficients being negative and should be avoided because the objective is to minimize the number of working periods. Setting ε to a value below zero, on the other hand, would penalize late processing and push the operations to as early in the day as possible and would not serve to shorten the working intervals.

4.2. *Solution framework*

In essence, we have a multi-criteria optimization problem. To develop an equipment schedule, we use a three-step approach related to sequential goal programming, along with a post-processor to assign operations to machines (Zhang and Bard, 2005a). The three steps correspond to solving three optimization problems associated with the previously defined objectives. The third problem involves a weighted sum of objectives (3a) and (3b).

The following parameters are used in the description of the approach:

θ_1 = a relaxation parameter for unprocessed mail, $0 \le \theta_1$;
θ_2 = a relaxation parameter for shifts, $0 \le \theta_2$;
θ_3 = a smoothing parameter, $0 \le \theta_3 \le 1$.

Sequential procedure

Step 1. (Model 1.) Solve the basic multi-level lot sizing model to minimize the ending inventory:

$$\psi_1 = \text{Minimize} \sum_{n\in N} v_n(48), \qquad (4)$$

subject to Equations (1a)–(1h)

Step 2. (Model 2.) Relax the volume of mail not processed in Step 1 for each operation n by the percentage θ_1.

Solve the following problem to obtain the minimum number of full-time shifts:

$$\psi_2 = \text{Minimize} \sum_{k\in K}\sum_{f\in F} \omega_{kf}, \tag{5a}$$

subject to Equations (1a)–(1h) and

$$v_n(48) \le \psi_{1n} + \theta_1 \rho_n, \quad \forall n \in N, \tag{5b}$$

$$\sum_{f:s_f\le t\le e_f} \omega_{kf} \ge \sum_{g\in G(k)}\sum_{n\in N(g)} r_{kg} Y_{ng}(t), \quad \forall t \in T,\ k \in K. \tag{5c}$$

In Equation (5b), ψ_{1n} is the ending inventory of operation n associated with the solution ψ_1 found in Step 1, and $\rho_n$is the volume of mail that can be processed at node n in one period (equivalent to the throughput of operation n).

Step 3. (Model 3.) Relax the number of shifts found in Step 2 by the percentage θ_2. Solve the following problem to minimize a combination of the number of startups and the weighted sum of working periods:

$$\psi_3 = \text{Minimize} \sum_{n\in N}\sum_{t\in T(n)}\sum_{g\in G(n)} \big[(1-\theta_3)Z^1_{ng}(t) + \theta_3(1-0.01t)Y_{ng}(t)\big]. \tag{6a}$$

subject to Equations (1a)–(1h), (5b)–(5c) and

$$\sum_{k\in K}\sum_{f\in F} \omega_{kf} \le \psi_2 + \theta_2\psi_2. \tag{6b}$$

Here, ψ_2 is the optimal number of shifts found in Step 2. Parameter θ_3 can be set in the range of zero and one. For values close to zero, the emphasis is on smoothing the schedule, that is, minimizing the number of times the machines switch operations or are turned on and off. When θ_3 is close to one, the emphasis is on compressing the schedule.

Notice in models 2 and 3, there is no integral requirement on the shift variables ω_{kf}. More formerly, we have the following result. The proof is given in Zhang (2003) and is based on the "consecutive one's property" of the corresponding **A** matrix.

Proposition 1. *The shift variables ω_{kf} will always be integral in an optimal solution to models 2 and 3 for r_{kg} and ψ_2 integer.*

Despite this result, our initial computational experience indicated that these models can present a formidable challenge for even the best commercial solvers. This can be explained in part by the following observation.

Proposition 2. *The mixed-integer program defined by model 1 is NP-hard in the strong sense.*

Proof. The basic idea is to consider a simplified version of the equipment scheduling problem whose objective is to minimize the volume of unprocessed mail subject to precedence constraints and show that a known NP-hard problem, in this case $F2|p_{ij} = 1, prec|\sum_i w_i U_i$, can be polynomially transformed into it. Brucker and Knust (1999) showed that the "easier" problem $F2|p_{ij} = 1, chains|\sum_i U_i$, is NP-hard in the strong sense. This implies that our problem also falls in this category. The details are given by Zhang and Bard (2005b). ■

Corollary 1. *The optimization problems associated with models 2 and 3 are NP-hard in the strong sense.*

5. Computational improvements

The proposed methodology involves solving three large-scale MILPs. To reduce the computational burden two major enhancements were developed. The first is a pre-processor designed to reduce the number of integer variables in the models; the second is a heuristic that uses the linear programming relaxation as a target and attempts to find integer solutions that are as close to it as possible.

5.1. *Pre-processing*

Notice that the only integer variables in the models are $Y_{ng}(t)$ whose number depends largely on the length of the working window in which an operation can be performed. By shortening the window of each operation, the total number of integer variables can be reduced. Whereas the end time of an operation is usually set according to the dispatching schedule and cannot be changed, the start time can be adjusted to reflect the fact that an operation will not be started until a certain accumulation of, say, 30 minutes of mail or until θ% of the total volume, has arrived.

To account for each of these cases, let $V_n(t)$ be the maximum volume of mail that can arrive at node n before the beginning of period t. When volume is measured by minutes of mail, the start time can be set to $t_1 = \min\{t : V_n(t) \ge \rho_n\}$, where ρ_n is the volume of mail that can be processed at node n in a 30-minute period. When volume is measured by a percentage θ of the total, the start time can be set to $t_2 = \min\{t : V_n(t) \ge \theta V_n(48)\}$, where $V_n(48)$ is the total volume associated with operation n during the day and θ is the percentage. We will refer to these two criteria as rule 1 and rule 2, respectively.

The algorithm now presented is used to calculate $V_n(t)$. In the description, $W_n(t)$ is the amount of mail processed in period t, and $A_n(t)$ is the amount of mail that arrives at node n in period t. The other notation was defined previously.

Step 1. Set $V_n(1) = v_n(0) = u_n$ (the mail left from the previous day), $W_n(0) = 0$, $t = 1$.

Step 2. While $t \le 48$, do

For each operation n, do the following:

Table 1. Results for different pre-processing rules

		Statistics of model 1			*Model 1*		*Model 2*	
Rules	*Start time*	*Number of rows*	*Number of columns*	*Number of integers*	*Solution*	*Time (seconds)*	*First IP solution*	*Time to first IP (seconds)*
0	No pre-processing	5346	6700	849	99 998	9	253	350
1	t_1	5079	6274	694	101 390	6	254	149
2	t_2	5077	6274	693	99 998	6	252	113
3	$\text{Max}\{t_1, t_2\}$	4900	5938	618	101 390	3	251	90
4	$\text{Min}\{t_1, t_2\}$	5251	6610	769	99 998	5	254	239

(i) compute $A_n(t) = \sum_{i\in I(n), t\in T(i)} a_i(t) + \sum_{p\in P(n)} f_{pn} W_p(t-1)$
(ii) compute $W_n(t) = \min\{v_n(t-1) + A_n(t), \sum_{g\in N(g)} m_g(t)\rho_n\}$, if $t \in T(n)$; $= 0$, otherwise
(iii) compute $V_n(t) = V_n(t-1) + A_n(t)$
(iv) compute $v_n(t) = v_n(t-1) + A_n(t) - W_n(t)$

End of operation n
$t = t + 1$
End of all periods

Step 3. Stop

The algorithm pushes as much mail through the system as possible in each period, limited only by the capacity of the machines. As a result, mail is sent to downstream operations as early as possible. The algorithm has complexity $O(|N| \times |T|)$, where $|N|$ and $|T|$ are the cardinalities of sets N and T, respectively.

To test the effectiveness of the pre-processing algorithm, each model was solved with Xpress (Dash Optimization, 2002) for the various combinations of start times. Table 1 presents the statistics for model 1 for a typical Monday. The number of variables and constraints only increases slightly for models 2 and 3 due to the addition of ending inventory constraints (5b), and the staff covering constraints (5c) and (6b), so the statistics for those models are not reported. The "number of integers" column refers to the number of integer variables in the models. The objective values and solution times are included for model 1, which runs very quickly. For model 2, we report the objective values associated with the first integer point found, and the solution times to reach that point. For rule 3, we set the start time to max $\{t_1, t_2\}$ and for rule 4, we set the start time to min $\{t_1, t_2\}$. The parameter θ was set to 10% in all cases.

As an aside, the large discrepancy in run times between the two models is due to a combination of their objective functions and feasible regions. In model 1, we are trying to process as much mail as possible, and because the amount of resources (machines) available is usually more than adequate, there are many feasible solutions that conform to this objective. In model 2, if we think of the staff as a resource, the feasible region is much tighter so minimizing the number of shifts turns out to be a much more difficult problem.

From Table 1, we see that the number of integer variables drops from 849 under no pre-processing to 618 under rule 3. Computation times decreased proportionately. For model 1 there was a 67% drop from 9 to 3 seconds, and for model 2, a 75% drop from 350 to 90 seconds in the time required to find the first integer solution. In most cases, the objective values remained about the same. For model 1, the amount of unprocessed mail increased slightly from 99 998 to 101 390 pieces when the 30-minute rule was used. Overall, rule 3 gave the best results and was used in our implementation.

5.2. *LP-based heuristic for model 2*

Even with pre-processing, model 2 can still be quite difficult to solve. The problem has to be solved seven times, one for each day of the week and depending on the instances, it could take anywhere from 1 or 2 minutes to 10 minutes to find an integer solution. However, for mixed integer programs that are not combinatorial in nature, it is often possible to construct high quality solutions from the LP relaxation. This observation provided the motivation for our LP-based heuristic.

To begin, we solve the LP relaxation associated with model 2 to obtain target values for the decision variables. This can be done in a few seconds. The results are used to construct a replacement IP for model 2 with the objective of minimizing the sum of the absolute deviations of the integer component of the solution from the target values. Specifically, let $Y_{ng}^{\text{LP}}(t)$ be the value of $Y_{ng}(t)$ in the LP relaxation of model 2, $d_{ng}(t)$ be the deviation from $Y_{ng}^{\text{LP}}(t)$, and define β to be a "small" weight associated with the original objective, the number of full time shifts. We now solve

$$\text{Minimize} \quad \sum_{t\in T}\sum_{n\in N}\sum_{g\in G(n)} d_{ng}(t) + \beta \sum_{k\in K}\sum_{f\in F} \omega_{kf}, \tag{7a}$$

subject to Equations (1a)–(1h) and (5b)–(5c)

$$d_{ng}(t) \geq Y_{ng}(t) - Y_{ng}^{\text{LP}}(t), \quad \forall g \in G(n), t \in T(n), n \in N, \tag{7b}$$

$$d_{ng}(t) \geq Y_{ng}^{\text{LP}}(t) - Y_{ng}(t), \quad \forall g \in G(n), t \in T(n), n \in N, \tag{7c}$$

$$d_{ng}(t) \geq 0, \quad \forall g \in G(n), t \in T(n), n \in N, \tag{7d}$$

Table 2. Results of the approximation algorithm for Monday

		LP based heuristic		*Original problem*	
Rules	*Start time*	*First IP solution*	*Time to first IP (seconds)*	*First IP solution*	*Time to first IP (seconds)*
0	No pre-processing	254	55	253	350
1	t_1	251	16	254	149
2	t_2	252	20	252	113
3	Max$\{t_1, t_2\}$	252	13	251	90
4	Min$\{t_1, t_2\}$	254	33	254	239

Constraints (7b) and (7c) together define $d_{ng}(t) = |Y_{ng}(t) - Y_{ng}^{\mathrm{LP}}(t)|$, the deviation of the IP solution from the relaxed LP solution of the original problem. The objective (7a) is to minimize a combination of the absolute deviation (term 1) and the original objective function (term 2). Here, term 1 acts as the primary objective function and term 2 as a secondary objective. Let us call Equations (7a)–(7d) the model 2 approximation problem.

Table 2 presents the computational results for both the model 2 approximation problem and the original problem under different pre-processing rules using the same data for Monday. The "Time to first IP" column refers to the time to get the first integer solution. This value is what is reported in the "First IP solution" column. For the model 2 approximation problem, the IP value is not Equation (7a) but the value of the model 2 objective function, Equation (5a). The analysis was conducted for the different pre-processing rules to determine the degree of uniformity of the results. As the data in the table show, solving the approximation problem yields similar or better solutions in only about 14% of the time required to solve the original problem. These results were consistent across the week.

The use of target solutions has been studied by others in different contexts (e.g., see Morton *et al.* (2003)). We believe that one of the reasons why convergence is often so rapid is because the search is indirectly limited to a small neighborhood of the target solution. That is, we are looking for a local rather than a global solution.

6. Computational results

A comprehensive set of experiments was conducted to evaluate the quality of the schedules produced by the three-stage methodology, as well as to develop some insights into the relationships between the equipment and staff schedules. All models were implemented using Mosel, the modeling language provided by Dash Optimization and solved with Xpress, their general linear and integer programming software (Dash Optimization, 2002). The computations were performed on a 1.13 GHz Pentium III PC with 256 MB of RAM.

Table 3. Results from model 1

Measure	*Total arrivals*	*Number of pieces processed*	*Number of pieces unprocessed*
Mail volume	5 047 628	4 947 630	99 998
Percentage (%)	100	98.2	1.8

The input data, in the form of letter volume arrival profiles and end-of-run reports, were provided by the Dallas P&DC and reflect a typical week of activity in the facility (accounting period 7, week 3, fiscal year 2002). The end-of-run reports are used to determine the branching ratios at each node in the network in Fig 2.

6.1. *Feasibility of the schedule*

Table 3 presents the total arrivals, the number of pieces processed, and the number of pieces remaining on a typical Monday. The results were obtained by solving model 1 with the objective of minimizing the ending inventory. Through the day, more than 5000 000 letters arrived at the facility and 98.2% of them were processed in the normal flow. The remaining 99 998 were held over in inventory until the next day.

The amount of unprocessed mail comprised 1.8% of the total arrivals. Of those mail pieces, it was found that 1% out of the 1.8% was mail that arrived too late to be completed, whereas the remaining 0.8% was back flow mail. That is, mail that must be returned to a previous operation and cannot be processed on the current day because the operation had been closed.

Model 1 consists of 4900 constraints and 5938 variables, 618 of which are integer valued. It was solved to optimality in 3 seconds. The stopping criterion was an absolute gap of 1000 mail pieces.

6.2. *Exact staffing cost and the importance of model 2 in the three-step approach*

The number of shifts indicated by the solution of model 2 only provides an approximation of the daily staffing cost. To determine the exact value imposed by the equipment schedule generated by our methodology, we used the weekly staff scheduler (SOS), an implementation of the model presented in Bard *et al*. (2003), for the computations. The overall procedure was as follows.

Step 1. Solve models 1, 2 and 3 sequentially for each day of the week, carrying over the ending inventory to the next day.

Step 2. Solve the weekly staff scheduling problem (SOS) using the solutions obtained from model 3 as input.

At Step 2, the settings for SOS were: full-time to part-time employee ratio of at least 4:1, lunch breaks, 2-hour start time windows for weekly schedules, and no requirement for 2 days off in a row.

Table 4. Result from SOS without including the surrogate constraint

Performance measure	*Almost pure smoothing* $\theta_3 = 0.01$	*More smoothing* $\theta_3 = 0.2$	*Less smoothing* $\theta_3 = 0.4$	*Less compression* $\theta_3 = 0.6$	*More compression* $\theta_3 = 0.8$	*Almost pure compression* $\theta_3 = 0.99$
Cost ($)	139 978	133 335	131 226	134 784	137 086	139 161
Full-time regulars	112	105	104	108	109	112
Part-time flexibles	28	26	26	27	27	28

We now establish the importance of model 2, which provides an indirect estimate of daily staffing cost by controlling the number of full-time shifts. As mentioned, constraint (6b) is included for this purpose. By varying the relaxation parameter θ_2, different schedules can be obtained. When θ_2 is set to a large value, say two, Equation (6b) will not be active in model 3. When θ_2 is set to a small value, say 0.1, its presence becomes decisive.

6.2.1. *Exact cost of running facility when θ_2 is set to a large value*

Table 4 presents the staffing cost from SOS as well as the number of full-time and part-time workers necessary to run the automation equipment when the surrogate constraint is omitted from the solution framework, i.e., when θ_2 is set to a large value, say two. From the table we see that the difference between the best and worst solutions is roughly 6.7%. Note that the slight cost difference between the two solutions with the same number of workers is due to the fact that a 1% optimality gap was used as the stopping criterion in SOS, and that part-time shifts are of different lengths and hence different costs.

In the analysis, various settings for the preference parameter θ_3 were tested in an effort to determine whether a smoothed schedule or a compressed schedule was better. From the table, we see that the best result is obtained by striking a balance between smoothing and compression when θ_3 is set to 0.4. For the output reported in Table 4, a time limit of 600 seconds was placed on each run. On average, the optimality gap was 2% at termination, but feasible solutions with a 4% gap were consistently obtained in about 18 seconds.

6.2.2. *Exact cost of running facility when θ_2 is set to a small value*

Table 5 presents the staffing cost and the number of full-time and part-time workers obtained from SOS when the surrogate constraint is set to be active, i.e., when θ_2 is set to a small value, say 0.1. This allows a 10% increase in the number of shifts above the minimum. Again, different settings for preference parameter θ_3 were evaluated.

Comparing the results in Tables 4 and 5, the best solutions found were \$120 935 and \$131 226, respectively, an impressive 9% difference. As we can see, the existence of constraint (6b) provides a link between equipment scheduling and staff scheduling, and is of critical importance to the quality of the solution.

When constraints (5b), (5c) and (6b) are included in model 3, the integer program becomes much harder to solve. The average gap was 6.17% for a time limit of 600 seconds. In addition, it takes roughly 60 seconds to get the first feasible solution, much longer than before. When θ_2 is set close to zero, the problem gets even more difficult to solve. To help relieve the computational burden, we developed a "piece-by-piece" linear programming heuristic which is fully described in Zhang and Bard (2005b).

We also investigated the impact shift sets of different sizes on staffing requirements. The analysis included a set of three shifts (a shift every 8 hours), a set of nine shifts (a shift every 2 hours), and a set of 18 shifts (a shift every hour), all starting between 7 a.m. and 11 p.m. The results suggested that the number of shifts does not strongly affect the final costs produced by SOS and so are not reported. In our implementation, we choose the nine-shift set, primarily because the corresponding start times are very close to those used in practice.

The final analysis involved the exploration of different settings for θ_2 and θ_3. By decreasing the relaxation parameter θ_2 from two to zero, better solutions resulted for all values of θ_3. Setting $\theta_2 = 2$ makes the surrogate constraint (6b) redundant, and produced the highest costs; setting θ_2 close to zero provided the best results. As θ_2 approaches zero, however, model 3 becomes exceedingly difficult to solve. In view of this, it was decided to set $\theta_2 = 0.05$ and $\theta_3 = 0.4$ in

Table 5. Result from SOS including surrogate constraint, $\theta_2 = 0.1$

Performance measure	*Almost pure smoothing* $\theta_3 = 0.01$	*More smoothing* $\theta_3 = 0.2$	*Less smoothing* $\theta_3 = 0.4$	*Less compression* $\theta_3 = 0.6$	*More compression* $\theta_3 = 0.8$	*Almost pure compression* $\theta_3 = 0.99$
Cost ($)	132 067	130 793	126 585	130 000	120 935	121 729
Full-time regulars	104	104	100	104	96	96
Part-time flexibles	26	26	25	26	24	24

the baseline model. These values strike a balance between lower staffing cost and fewer startups, which is attractive from an operational point of view.

6.3. *Summary of computations*

The computational results for models 1, 2 and 3 are summarized in Table 6 for each day of the week. For model 1, the objective represents the amount of unprocessed mail (in thousands) remaining at the end of the day. The "%" column records the percent difference between the total processed mail and the total arrivals. Note that the large percentage of unprocessed mail (41%) on Saturday is due to the fact that delivery point sequencings is not performed on that day since there is no delivery on Sunday. The mail for these operations is carried over and processed on Sunday. Here, the negative percentage (−65%) indicates that the total volume processed is greater than the total arrivals due to the carryover.

For model 2, we present the results obtained by solving the approximation problem and stopping at the first integer solution. The "Root node obj" column indicates the relaxed LP solution of the original model, and the "Gap %" denotes the difference between the LP solution and the corresponding objective value of the approximation problem. For model 3, a time limit of 600 seconds was imposed on the computations because little improvement was ever realized after this amount of time. The "Root node obj" column indicates the LP solution, and the "Gap %" denotes the difference between the LP solution and the best IP solution found. The best bound refers to the lower bound at the end of the computations. The number of working periods, startups, and shifts are also presented. The average reduction in startups per day between models 2 and 3 is 25.7% and the average reduction in working periods is 2.2%.

In general, the schedules provided by model 3 exhibit continuous operations with few setups and start-stop patterns. They also guarantee the prompt processing of mail, and as discussed presently, impose much lower labor costs than the schedules currently being used.

7. Quality of schedule

To validate the methodology, we ran our equipment scheduler using volume arrival data provided by the Dallas P&DC for a typical week, and asked an in-plant support specialist to evaluate the results for the automation operations. To better match the actual processed volume in this week, it was necessary to increase the input volume to our scheduler by 10%. We ran SOS and compared the quality and cost of our solution with the solution obtained from SOS using an equipment schedule (obtained with SiteMeta) in use by the facility at the time of the study.

A key observation concerns the staffing cost associated with the equipment schedule generated by our three-stage methodology and how it compares with the SiteMeta schedule. Table 7 gives the staffing cost as well as the number of Full-Time Regulars (FTRs) and Part-Time Flexibles (PTFs) needed in either case. As can be seen, our schedule calls for 112 FTRs and 28 PTFs compared to the SiteMeta schedule which calls for 162 FTRs and 40 PTFs.

The higher cost associated with the SiteMeta schedule, a striking 44%, is in part due to the adjustments that the supervisors made to account for uniformity of start times, personal preferences for work assignments, and overestimates in the number of periods required for an operation. Whereas some of these considerations should be taken into account as the model is refined, it is self-defeating to adhere to business-as-usual policies when there is no compelling reason to do so. Adding restrictions simply increases the staffing cost. To improve the usefulness of the results, however, it will be necessary to address several points, such as weekly arrival fluctuations, inaccuracies in the branching probabilities, and the need to further reduce computation times.

8. Discussion and conclusions

The general scheduling problem in high volume facilities, such as mail processing and distribution centers, is too complex to handle in a single model. A hierarchical approach is often the most effective way to deal with the interplay among equipment, workers, and operational restrictions. In this paper, we examined the equipment scheduling aspect of the problem with the implicit goal of minimizing the labor costs associated with the permanent workforce. However, labor costs are only one of several criteria that are used to judge the quality and acceptability of a schedule. In our case, service standards as measured by the percentage of arriving mail processed per day, and the number of machine startups, were the two additional criteria used in the analysis.

To address each of these concerns, the equipment scheduling problem was formulated as a large-scale MILP and solved sequentially using a three-stage methodology. An efficient pre-processing heuristic was developed to reduce the number of variables, and an approximation scheme resembling a LP rounding heuristic was developed to speed a major portion of the computations. Together, they reduced solution times by more than 90% without any visible sacrifice in solution quality.

In fact, the results indicate that implementing the three-stage methodology will lead to substantial savings nationwide and a greater understanding of facility operations. Optimizing the equipment schedule is the most direct way of achieving cost reductions while providing managers with a valuable tool for increasing the efficiency of P&DC operations. If only 50% of the benefits are realized, roughly \$1600 000 annual savings per facility can be achieved for the letter section alone.

Table 6. Computational results for models 1, 2 and 3

	Model 1				*Model 2: $\theta_2 = 0.01$*					*Model 3: $\theta_2 = 0.05, \theta_3 = 0.4$*						
Day	*Objective function (×1000)*	*Root node obj*	*Time (seconds)*	*%*	*Obj (No. shifts)*	*Root node obj*	*Gap %*	*Number of start-ups*	*Number of periods*	*Objective function*	*Root node obj*	*Best bound*	*Gap %*	*Number of start-ups*	*Number of periods*	*Number of shifts*
Monday	99.98	99.98	3	1.8	252	236.70	6.46	388	3011	941.82	909.10	918.66	5.2	290	2925	265
Tuesday	94.98	94.98	4	1.1	236	220.61	6.98	392	2633	935.28	862.10	870.00	7.5	269	2617	248
Wednesday	104.46	104.46	5	0.3	233	216.30	7.72	357	2717	908.60	856.12	861.26	5.5	258	2636	245
Thursday	158.83	158.75	5	1.1	240	224.68	6.82	339	2884	963.88	909.68	915.54	5.3	265	2829	253
Friday	207.76	207.76	5	0.9	239	221.28	8.01	372	2841	918.86	866.30	868.70	5.8	262	2772	251
Saturday	1651.91	1651.91	2	41.0	143	125.01	14.40	229	1911	575.32	522.76	529.30	8.7	164	1855	151
Sunday	53.72	53.50	2	−65.0	156	147.28	5.92	218	1743	543.48	512.86	516.40	5.2	183	1708	164

Table 7. Comparison of staffing costs

Performance measure	*Three-stage methodology*	*SiteMeta schedule*
Cost	\$144 544	\$208 221
FTRs	112	162
PTFs	28	40

The proposed methodology is being implemented in a decision-support system that is currently running at several P&DCs. One interesting extension would be the development of a short-term scheduling model that addresses the weekly fluctuations in mail arrivals and employee absenteeism. Such a model would provide a valuable tool for the efficient day-to-day management of the facility.

Acknowledgment

This work was supported in part by the National Science Foundation under grant # DMI-0218701.

References

Bard, J.F., deSilva, A.H., Feo, T.A. and Wert, S.D. (1993) Design of semi-automated mail processing facilities. *IIE Transactions on Design & Manufacturing*, **25**(4), 88–101.

Bard, J.F., Binici, C. and deSilva, A.H. (2003) Staff scheduling at the United States Postal Service. *Computers & Operations Research*, **30**(5), 745–771.

Beaumont, N. (1997) Scheduling staff using mixed integer programming. *European Journal of Operational Research*, **98**(3), 473–484.

Berman, O., Larson, R.C. and Pinker, E. (1997) Scheduling workforce and workflow in a high volume factory. *Management Science*, **43**(2), 158–172.

Brucker, P. and Knust, S. (1999) Complexity results for single-machine problems with positive finish-start time-lags. *Computing*, **63**, 299–316.

Brusco, M. and Jacobs, L.W. (1998) Personnel tour scheduling when starting time restrictions are present. *Management Science*, **44**(4), 534–547.

Dash Optimization (2002) *Xpress-Mosel: User Guide*, Dash Optimization, Englewood Cliffs, NJ.

Jarrah, A.I.Z., Bard, J.F. and deSilva, A.H. (1994a) Solving large-scale tour scheduling problems. *Management Science*, **40**(9), 1124–1144.

Jarrah, A.I.Z., Bard, J.F. and deSilva, A.H. (1994b) Equipment selection and machine scheduling in general mail facilities. *Management Science*, **40**(8), 1049–1068.

Malhotra, M.K., Ritzman, L.P., Benton, W.C. and Leong, G.K. (1992) A model for scheduling postal distribution employees. *European Journal of Operational Research*, **58**, 374–385.

Morton, D.P., Popova, E., Popova, I. and Zhong, M. (2003) Optimizing benchmark-based utility function, Working Paper, Graduate Program in Operations Research & Industrial Engineering, University of Texas, Austin, TX.

Zhang, X. (2003) Model and solution techniques for machine scheduling problems in high volume factories. Ph.D. dissertation, Graduate Program in Operations Research & Industrial Engineering, The University of Texas, Austin, TX.

Zhang, X. and Bard, J.F. (2005a) A bi-criteria multi-period machine assignment problem. *European Journal of Operational Research*, to appear.

Zhang, X. and Bard, J.F. (2005b) Comparative approaches to equipment scheduling in high volume factories. *Computer and Operations Research*, to appear.

Biographies

Xinhui Zhang is an Assistant Professor of Industrial Engineering at Wright State University. He received his Ph.D. in Operations Research & Industrial Engineering from the University of Texas at Austin. His research interests are in mathematical programming and optimization, especially solving large-scale problems that arise in manufacturing, logistics and transportation, media management, service operations, and engineering design. He has participated in and led several research projects involving airline crew recovery, advertising allocation for television networks, and equipment scheduling for mail processing facilities. He is an active member of INFORMS. Prior to beginning his career in operations research, he was a mechanical engineer specializing in CAD/CAM and robotics optimization, and won the outstanding award (No. 2) in the first national students' software competition in China.

Jonathan F. Bard is a Professor of Operations Research & Industrial Engineering in the Mechanical Engineering Department at the University of Texas at Austin. He holds the Industrial Properties Corporation Endowed Faculty Fellowship, and serves as the Associate Director of the Center for the Management of Operations and Logistics and as the Area Coordinator for the OR&IE Graduate Program. He received a D.Sc. in Operations Research from The George Washington University. He has previously taught at the University of California—Berkeley and Northeastern University. His research interests are in airline operations, the design and analysis of manufacturing systems, postal operations, and vehicle routing. Prior to beginning his academic career, he worked as a program manager for the Aerospace Corporation and as a systems engineer for Booz, Allen & Hamilton. He is currently the Editor of *IIE Transactions on Operations Engineering* and serves on the Editorial Board of several other journals. He is a fellow of IIE and a senior member of IEEE and INFORMS. In the past, he has held a number of offices in each of these organizations. His research has been published in a wide variety of technical journals.

Contributed by the Location and Transportation Modeling Department