

SEPARATING LOGISTICS FLOWS IN THE CHICAGO PUBLIC SCHOOL SYSTEM

DONALD D. EISENSTEIN and ANANTH V. IYER

University of Chicago, Chicago, Illinois

(Received July 1994; revisions received January, June, July 1995; accepted August 1995)

The central Chicago Public School warehouse was responsible for the distribution of supplies to 600 schools, including over \$10 million annually of engineering and educational supplies. The system was fraught with problems—deliveries were not made according to schedule, schools were hoarding inventories, and some schools were paying a premium for reliable service from third party suppliers. This paper reports how we improved this logistics system. We built a mathematical model of the system, validated our model using historical data, and used the model to evaluate the impact of potential changes to the system. Our recommended changes were implemented throughout the system. We report the impact on system performance. The redesigned system shows a dramatic reduction in lead times, a reduction in capacity requirements, and an overall reduction in system costs.

The Chicago Public School (CPS) System consists of 600 schools to which over \$100 million of food, engineering, and educational supplies are distributed annually. At the start of our study in 1992, this distribution was managed by a central warehouse, which also provided daily mail service among the schools and school board offices. We examined the distribution of engineering and educational (E&E) supplies. These items, which constituted \$10 million in annual sales volume, ranged from small boxes of paper clips and pencils to larger items such as copy paper, floor cleaner, and rock salt.

By December 1992 the performance of the E&E distribution system was deteriorating. School principals complained that deliveries were unreliable and lead times were long. To compensate, schools either carried large safety stocks of inventory or relied on expensive third party suppliers. The Financial Resource Advisory Committee, an independent not-for-profit agency, contacted us to research these distribution system problems and recommend changes. This paper documents that effort.

Our approach was to: collect and analyze system data, build models to evaluate the impact of design changes, implement our recommendations, and report on the observed changes in system performance. When the recommended changes were implemented, on-time delivery increased to over 98% from 45%, lead times for the large majority of items decreased from over two weeks to two days, and required truck capacity decreased by 40%. As a result of these changes, the warehouse director reported a decrease in system costs of \$300,000 annually.

Our main design change was to separate the E&E distribution system into two components, each with its own set

of items, resources, and delivery schedules. We term each such independent distribution system a *logistics system*. We modeled the *logistics interactions* in the original system. These logistics interactions caused a loss of capacity and large lead times. We decreased these logistics interactions by replacing the original system for E&E distribution by two separate logistics systems.

1. RELATED WORK

Recent literature in applied logistics describes the benefits of tailored logistics systems. Fuller, O'Conner and Rawlinson (1993) discuss how aggregating logistics systems across many products or customers forces an averaging of costs and prevents efficiencies that can be generated by customizing logistics systems by customer or product type. Byrnes and Shapiro (1991) discuss an industrial products manufacturer who dramatically decreased logistics costs by separating logistics flows into two systems. The first system used planned deliveries in a standing order arrangement for large distributors, while the second used a traditional ship-to-order delivery system for smaller distributors. A Harvard Case (1990) discusses the separation of logistics flows through the use of a quantity discount tied to an order volume and a lead time. Large distributors would pay a lower unit cost, get deliveries directly from the central warehouse in large volumes, and face a longer lead time. Smaller distributors would get overnight delivery, but pay a higher cost per unit. Smith and Barry (1991) describe a pharmacy system where satellite pharmacies were used to provide small doses of medicines quickly while a central pharmacy focused on larger doses. The central pharmacy

Subject classifications: Inventory/production: policies review/lead times. Programming: integer applications. Transportation: vehicle routing.
Area of review: OR PRACTICE.

had a larger lead time, but focused on maximizing dose picking efficiency. These examples illustrate different logistics systems and their associated lead times.

Our focus was on modeling the impact of separating logistics flows into logistics systems by modeling the impact of logistics interactions. We built a model of the lead time generated by the system as a result of the interaction between order volume and truck capacity allocation. This is in the same spirit as queueing models of lead time used in the analysis of inventory systems (see Zipkin 1986, and Karmarkar 1987). Our model of lead time used the observed order arrival streams and deterministic packing models based on data from the CPS warehouse.

2. THE DESIGN AND OPERATION OF THE CHICAGO PUBLIC SCHOOL LOGISTICS SYSTEM

We describe the CPS warehouse operation, the sequence of decisions, and the rules used by the associated decision makers. We examined decisions made at three levels: *strategic*, *tactical*, and *operational*. The *strategic* level decision created logistics systems by partitioning items into groups with each group assigned its own trucking capacity. At the next level, *tactical* decisions assigned schools to routes, sequenced truck stops, managed the order taking system, located items in the warehouse, and coordinated the picking and packing of the items. Finally, *operational* decisions detailed the packing of the items on the trucks and handled last-minute rerouting and rescheduling. The following summarizes the main features of the system.

1. At the *strategic* level, the *Director* of the Department of Warehousing and Distribution determined how to allocate trucks between the distribution of food, engineering supplies, educational supplies, and mail delivery. He designed his system to send trucks as full as possible. He used separate sets of trucks for the distribution of food, for the delivery of mail, and for the delivery of all E&E supplies.
2. At the *tactical* level, the *Traffic Manager* routed the allocated trucks. Each day five truck routes were scheduled for E&E distribution. These "planned" routes were developed to keep the truck and driver highly utilized. The workload, consisting of load pickup at the central warehouse, driving time, and delivery time to schools, was designed to be an eight hour work day. The Traffic Manager planned the routes so that each school would be visited every two weeks (10 work days). With 600 total schools, there averaged 60 schools for delivery each day and therefore about 12 schools on each of the five daily routes. Since high schools are usually surrounded by neighboring elementary schools, truck routes typically included about two or three high schools and about nine or ten elementary schools.
3. The *Order Entry Clerks*, also at the *tactical* level, assigned each school a designated day, once every two weeks, to phone in its orders. For each order called in

by a school, the order entry clerks generated a pick-and-pack order list after confirming that the items were available and the school's budget was adequate.

One of the Order Entry Clerks handled the key engineering items. These were items that had large physical volume as well as high sales volume. At the start of each year, each school engineer provided monthly forecasts for the entire year of the requirements for each key item. The Order Entry Clerk entered these forecasts, just like phone orders, into the system each month throughout the year.

4. The *Warehouse Manager*, also at the *tactical* level, was in charge of picking items and moving them to the loading area. The Warehouse Manager supported a policy of *order integrity* for each school; that is, the entire outstanding order for a school should be delivered together. Therefore, the Warehouse Manager consolidated the entire order for a school at the loading area. This policy of order integrity simplified the bookkeeping for both the warehouse and the schools. The Warehouse Manager used forklifts to pick the bulky items and a shopping cart to pick the smaller items. The smaller items were collected and packed into small "make-up" boxes of various sizes. Each make-up box typically contained two to three items.
5. At the *operational* level, the *Warehouse Foreman* packed the truck. The foreman can create no more than 16 pallet loads to tile onto the bed of the truck and he must keep the load for a school consolidated on the truck so it can be easily unloaded.

Our first visit to the warehouse illustrated the foreman's problem. A typical load for a school consisted of large bags of rock salt, large buckets of floor cleaner, and neatly stacked boxes of copy paper. Stacked on top of these larger items were small make-up boxes of many different sizes. The foreman had to pack these loads onto 16 pallets while maintaining order integrity and order consolidation for each school. In addition, the foreman often faced loads for a route that exceeded truck capacity. He then had to decide which schools got delivery on such days and which would be delayed.

3. DATA COLLECTION AND ANALYSIS

Data were collected to estimate the lead times observed by schools, the volume of orders faced by the warehouse over time, the capacity requirements of each truck route, and the volume of orders placed to third party suppliers. The data sources included school surveys, school visits, and computer databases.

3.1. Lead Time

Phone surveys of school principals showed consistent complaints of late and erratic deliveries. To estimate the magnitude of this problem, we calculated the actual lead time faced by schools. This lead time information had to be obtained from hard copies of the delivery logs. We chose a




Figure 1. Histogram of observed delivery lead times for a sample of 30 schools for the original system.

sample of 30 schools and tracked all orders placed between August 1992 and November 1992. Figure 1 shows a histogram of these lead times. The mean lead time observed was 20.35 days with a standard deviation of 12.20 days. The observed 45% probability of delivery within two weeks indicated a basis for school complaints.

3.2. Physical Order Volume

We had data regarding the day each order was placed, the school placing the order, and the quantity of each item ordered. However, the physical volume of each item was not available in the database. We thus had to estimate the physical volume of school orders received by the warehouse. Our approach was to separate items into two classes.

- Class A—items of large physical volume and high sales volume (includes most engineering items and some educational items).
- Class B—items not in Class A, typically of small physical volume (predominately educational items, but including a few engineering items of large physical volume and low sales volume).

We then collected the precise physical volumes for each item in Class A, while for Class B items we formed an aggregate estimate of their volume.

We selected the items in Class A by examining the sales volume of all the items. Of the 1,800 items, just 21 items accounted for 50% of the sales volume. These 21 items were used by most schools and, as confirmed by warehouse personnel, were among the largest in terms of physical volume. These items included copy paper, floor cleaners, and rock salt, and each unit was typically a cubic foot or more in physical volume. These 21 items constituted Class A. Figure 2 shows the total volume of Class A items ordered over time.

The remaining 1,779 items constituted Class B. To determine an average volume per unit of a Class B item, we collected data on the average make-up box sizes and the number of items packed into each make-up box. This




Figure 2. Physical volume (cubic feet) of Class A and Class B items ordered over time. Note that the volumes of Class A items are significantly larger than Class B items.

analysis estimated that the volume contributed by each unit of Class B averaged 0.1 ft³. Figure 2 shows the estimated volume of Class B ordered over time.

We estimated that Class B items accounted for about 23% of the total physical volume and Class A items accounted for the remaining 77%. The total physical volume of orders over the five month period was estimated at 158,960 ft³.

3.3. Utilization of Trucks

The E&E trucks used by the CPS had an effective truck capacity of 700 ft³. Thus, the 158,960 ft³ of order volume over the five month period corresponds to 227 truck loads of physical volume. The truck logs over this period indicated that 375 trucks were actually used, indicating an average truck volume utilization of about 60%.

We examined the volume scheduled to be loaded on each route each day over the five month period. Even though the average truck utilization was 60%, over half of the routes exceeded the truck capacity at least once over the five month period. In such cases, the delivery of some orders was delayed. Overall, the truck routes were not well balanced in terms of physical load.

We then examined the routing of these trucks to determine if the routes were balanced in terms of travel and delivery time. We modeled the travel speed and pick-up times based on data collected from the driver's delivery logs. The model used Euclidean distances between locations; an average speed of 15 mph to compensate for the congestion of city roads and non-Euclidean travel; three hours to unload a full truck plus a fixed time of five minutes per stop; and 1.5 hours to load and pack the truck each morning. Across all 50 truck routes, our model estimated a mean round-trip time of 7.56 hours and a standard deviation of only 0.5 hours. This indicated that the routes were balanced in terms of travel and delivery time.

3.4. Third Party Suppliers

The CPS system permitted schools to order directly from third party suppliers. (The items ordered from third party suppliers were almost exclusively Class B items.) The typical third party supplier delivered within one or two days, but the cost averaged about 15% higher than the cost from the warehouse. The lower lead time was the main incentive for schools to order from third party suppliers. Thus, as the warehouse deliveries became erratic, the use of third party suppliers grew. There was insufficient data to track the details of the E&E items ordered from third party vendors, but it was estimated to be about \$1.5 million annually or 15% of a school's total budget for E&E supplies.

4. LOGISTICS INTERACTIONS

This section describes three types of logistics interactions. These interactions provide a link between underlying demand at schools, orders received by the warehouse, and lead times faced by schools. In the following sections, we provide a mathematical model of these interactions and show that they explain the poor on-time delivery and low capacity utilization of the system.

Combining volumes of many items typically generates an aggregate pooled volume with a lower coefficient of variation. This was the motivation of the warehouse director in specifying that the 1,800 E&E items share a common pool of resources, and thus form a single *logistics system*. The interactions we describe are of interest because they provide a countervailing effect to pooling.

4.1. Capacity Interaction Effects—A Description of Truck Packing

In subsection 3.3 we reported only a 60% capacity utilization for the trucks; however, in subsection 3.1, we reported on-time delivery of under 50%. We thus examined the daily volume of Class A and Class B orders (as shown in Figure 2) and its relation to daily truck capacity to help explain this apparent inconsistency.

Figure 3 shows a histogram of the combined volumes of Class A and Class B items ordered each day. The orders exceeded the available truck capacity of $5 \times 700 = 3,500 \text{ ft}^3$ on only 15% of the days. Thus, Figure 3 does not explain the low on-time delivery performance of the system. We suggest that there may be more to truck packing than this simple *volume* description.

We modeled a loss in trucking capacity due to the mixing of Class A and Class B items—a *capacity interaction* effect. An order for a school typically consists of large boxes of Class A items plus small make-up boxes of Class B items. To simplify unloading at the school, these items are consolidated on the truck. Thus, a typical delivery for a school was stacked on a pallet with large Class A items on the bottom and the small make-up boxes of Class B items on top. Another school can rarely share this pallet because the Class A items of the next school could not stack on top




Figure 3. Histogram of the total volume ordered each day. The daily capacity is $3,500 \text{ ft}^3$.

of the fragile and uneven make-up boxes of Class B items of the first school. Thus, when both Class A and Class B items are ordered by a school, no more than one school's order can be loaded on a pallet. The truck capacity of 700 ft^3 was partitioned into 16 pallets of volume, and therefore, a school with a load of W will occupy approximately $\lceil W/(700/16) \rceil$ pallets.

Figure 4 examines the volume of Class A and Class B orders in terms of the *pallet loads* of volume generated each day. Figure 4 shows that on 33% of the days the number of pallets required exceeds the available capacity of $16 \times 5 = 80$ pallets. Thus, consideration of capacity interaction effects indicates a larger delivery lead time that is more consistent with the low on-time delivery reported in subsection 3.1.

4.2. Feedback Interaction Effects—A Description of Class B Variance

Figure 2 shows the fluctuation in the aggregate volume of Class B orders over time (the coefficient of variation was 0.41); however, surveys and interviews with school clerks suggested that the underlying demand volume of Class B items was nearly constant over time for each school. We




Figure 4. Histogram of the total pallet loads ordered each day. The daily capacity is 80 pallets.

will provide a model of *feedback interaction* effects that will explain this variation in Class B orders faced by the warehouse when the underlying demand for Class B items at the schools was considered to be relatively smooth.

We develop a model where the erratic orders in Figure 2 are a *consequence* of the interaction between school inventory policy and delivery lead time. In our model, the schools follow a simple *one-order-outstanding* inventory policy. This policy implies that a school will not place an order if another order placed by that school is outstanding at the warehouse. Thus, each school orders enough to cover the demand experienced since its last order. A school following a one-order-outstanding policy must keep large inventories or rely on third party suppliers so that erratic and long lead times do not result in stockouts.

The impact of a one-order-outstanding policy followed by a school on the variance of orders received by the warehouse can be understood by a simple model of a random sum of random variables (see Ross 1980). Let the lead time, L , observed by schools have a mean m_L and a variance σ_L^2 . Let the underlying demand volume in period i be D_i with mean m_D and variance σ_D^2 . If X is the physical volume of the order received by the warehouse for an individual school, then, under a one-order-outstanding policy, we have $X = \sum_{i=1}^L D_i$. The mean and variance of the order size is $\text{Mean}(X) = m_D m_L$ and $\text{Var}(X) = (m_L \sigma_D^2) + (m_D^2 \sigma_L^2)$. This expression shows that uncertain lead times accentuate the effect of the order size variance. This indicates a feedback effect—erratic delivery lead times faced by a school cause a large order variance to the warehouse which, in turn, causes erratic delivery lead times to the schools.

If delivery lead times to schools were truly independent, then the resulting variance in school order volume under this model would be much smaller than that observed in Figure 2 (due to a substantial pooling effect). However, the deliveries to schools on the same route are correlated. If a particular set of schools receive delivery on a route, then the other schools on the route are less likely to receive delivery due to capacity limitations. For instance, it was common for a set of large high schools to dominate the load on a truck in a given period and the rest of the schools to dominate delivery during the next period. We will show that our one-order-outstanding model of school behavior closely reflects the variance of Class B orders observed in the system, once the correlation between school delivery lead times is accounted for.

4.3. Order Entry Interaction Effects—A Description of Class A Variance

Figure 2 shows the volume of orders for Class A items over time. A distinct pattern of two weeks of high volume followed by two weeks of low volume was observed. We provide a model of *order entry interaction* to explain this variance.

The Class A items correspond closely to the key items, whose demands are forecasted by the engineer. Typically,

the order clerk would spend the first two weeks of each month entering the forecasts for that month into the order entry system. Thus, if a school was routed on the first and third Tuesday of each month, the clerk would schedule the *entire* monthly forecast of engineering items for delivery on the *first* Tuesday of the month. The rest of the warehouse personnel could not distinguish between the forecasted orders and actual orders called in by the schools. This would, in effect, put the warehouse in a bind at the beginning of every month. Due to the policy of order integrity, these large surges in engineering orders would result in some schools not receiving delivery on their scheduled day. Therefore, there was an interaction between the order entry for Class A items and the delivery lead times for Class B items.

5. THE TACTICAL MODEL

In this section:

1. We provide a mathematical model of the impact on lead time of the three interaction effects discussed in the last section. In Section 6 we will use this model to compare the model-generated lead time distribution with the observed lead time distribution in Figure 1. This will provide validation for our model.
2. We model the impact of changes to the original system. These changes include the elimination of capacity interaction effects and the elimination of order entry interaction effects. We then provide an algorithm to generate new truck routes.

5.1. Modeling the Original System

The time period t in the model corresponds to a two week delivery cycle. Each school year is naturally divided into two segments; August through December and January through May. We model only the first segment, and we divide it into 10 time periods, i.e., $t = 1, 2, \dots, 10$. In each time period, each school s is assigned to one truck route. For the original system there are 50 routes that repeat for each time period, i.e., $r = 1, 2, \dots, 50$. Let $S(r)$ denote the set of schools allocated to route r .

At the start of period t , the foreman is faced with a Class A outstanding load (A_{st}) and a Class B outstanding load (B_{st}) for each school s . For each route r , the foreman must then choose from among the schools $s \in S(r)$ those schools that will receive delivery in that period. The foreman maintains order integrity by delivering $A_{st} + B_{st}$ if school s receives delivery in period t . The foreman attempts to maximize truck capacity utilization (the warehouse director's objective) while satisfying constraints related to truck capacity and Class B picking capacity. The solution of this two-dimensional packing problem determines the schools that receive delivery each time period.

The delivery decisions made by the foreman in time t affect $A_{s,t+1}$ and $B_{s,t+1}$. However, we must model the foreman's delivery decision myopically, because he uses only information regarding the loads available at the beginning

of time t to make delivery decisions in time t , without any knowledge of the order stream he will face in period $t + 1$.

For each of the 10 periods, we solve a packing problem for each of the 50 routes. We thus solve 500 integer programs, each of which is a two-dimensional bin packing problem.

The following integer program models the foreman's packing decision for a particular time period t and particular route r .

$$\begin{aligned} & \text{Maximize} && \sum_{s \in S(r)} c_{st} x_{st} \\ & \text{subject to} && \sum_{s \in S(r)} p_{st} x_{st} \leq 16, \\ & && \sum_{s \in S(r)} u_{st} x_{st} \leq 96, \\ & && x_{st} \in \{0, 1\}, \quad \text{all } s \in S(r). \end{aligned}$$

The binary decision variables, x_{st} , are 1 if school s gets delivery in time period t , and 0 otherwise. The warehouse director had the objective of sending out trucks as full as possible every period. Thus, the objective function maximized by the foreman in each period for each route is to maximize the load on the truck. Thus, the coefficient c_{st} is set to $A_{st} + B_{st}$, where A_{st} is the Class A volume load outstanding, and B_{st} is the Class B volume load outstanding for school s at time t .

The first constraint enforces the limit of 16 available pallets for each truck, where p_{st} is the number of pallet loads required by school s in time period t . This load in pallets is expressed as $p_{st} = \lceil (A_{st} + B_{st}) / (700/16) \rceil$. While this expression only approximates the packing of a truck, the fact that we round up to the nearest integer makes this approximation reasonable.

The picking of items was done the day before shipment. The Class A items are easily picked by a forklift and so do not pose a constraint to the system, but Class B items are much more time consuming to pick. Dividing the picking day between five trucks left about 96 minutes to pick Class B orders for each truck. Furthermore, the time to pick each Class B item was estimated to average 0.685 minutes, so $u_{st} = 0.685 B_{st}$. The second constraint therefore states that a maximum of 96 minutes is spent picking the Class B items for a particular truck.

Modeling Class A and Class B Loads. We now model the loads A_{st} and B_{st} and describe the algorithm used to update their values each time period. These updated values of A_{st} and B_{st} are used to form the objective and constraints of the integer programs for each time period.

Class A orders are generated from the forecasts of demand provided by school engineers for 1992. We denote f_{st} as the forecast for Class A items for school s in period t . Since the forecasts are monthly, while our time periods are two-week intervals, we assign the entire forecast to the first order period of the month to model the batched order entry of Class A items. Thus, f_{st} is the entire monthly forecast for t odd and $f_{st} = 0$ for t even. The order for

Class A items for school s entered into the system each period t is denoted by e_{st} . Thus, the outstanding Class A order for school s in period t is $A_{st} = \sum_{k=L_{st}}^{t-1} e_{sk}$, where $e_{st} = f_{st}$ and L_{st} is the last period in which school s received delivery before period t .

Class B orders are modeled by assuming an underlying steady demand for Class B items at each school set to the actual average historical demand, and we let UB_{st} be this smoothed underlying demand for Class B items for school s in period t . The B_{st} values are then generated according to a one-order-outstanding policy. The model assumes that 30% of the schools would use a third party vendor if the warehouse was late on a shipment, based on data collected from warehouse personnel. We set $H_s = 0$ if school s never uses a third party supplier and set $H_s = 1$ if it might.

If there was no delivery to school s in period $t - 1$, then no more orders are received from school s because of the one-order-outstanding policy; that is, if $x_{s,t-1} = 0$ then $B_{st} = B_{s,t-1}$. However, if school s did receive a delivery in period $t - 1$, then if school s does not place orders from the third party supplier, the order outstanding is the total Class B demand that has occurred since the delivery before period $t - 1$; that is, if $x_{s,t-1} = 1$ and $H_s = 0$, then $B_{st} = \sum_{k=1+L_{s,t-1}}^{t-1} UB_{sk}$. If school s does place orders from a third party supplier and did receive a delivery in period $t - 1$, then the Class B order outstanding is merely the demand during period $t - 1$; that is, if $x_{s,t-1} = 1$ and $H_s = 1$, then $B_{st} = UB_{s,t-1}$.

Modeling Lead Times. The iterative solution of the integer programs determines the schools that receive delivery each period. At each period t we record lead times for delivery of Class A and Class B orders for each school s for which $x_{st} = 1$ as follows: For all periods $j = L_{st}$ to $t - 1$, record the lead time for each Class A order placed in period j as $t - j$ periods. Also, for each Class B order placed in period L_{st} , record the lead time for each order as $t - L_{st}$. Also update $L_{sk} = t$ for each period $k = t + 1, \dots, 10$.

We execute this procedure for all $t = 1, 2, \dots, 10$ and thus generate lead times estimates. In Section 6 we will show that the model-generated lead time distribution closely matches the observed distribution.

5.2. Modeling Changes to the System

In this section, we provide details of changes to the model that we make to study the impact of separating logistics flows, smoothing Class A forecasts, and rerouting of trucks. We first model the impact of separating the Class A and Class B logistics systems. To model the larger Class A logistics system we modify the parameters and algorithms of the original system model. The model for the Class A logistics system first sets all $B_{st} = 0$, thus making the objective function $c_{st} = A_{st}$. The capacity interaction effect has been removed and thus $p_{st} = A_{st}$, because the Class A items of one school can be stacked on the Class A items of

another. In addition, the Class B picking constraint is removed. The Class B logistics system is much smaller in volume and is therefore not modeled explicitly. (The Class B logistics system implementation will be discussed in the next section.)

To model the smoothing of Class A forecasts, and thus elimination of order entry interaction effects, the monthly engineering Class A forecast is divided evenly between both scheduled delivery days in the month, so that $e_{st} = (f_{s,t} + f_{s,t+1})/2$ for each t odd and $e_{st} = (f_{s,t-1} + f_{s,t})/2$ for each t even.

Routing changes are modeled by changing the partition of schools into route sets $S(r)$, while maintaining that each school s is assigned exactly one route r per time period. In some of the model runs, the traffic manager generated routes, in other runs we generated a set of routes. To generate these new routes we checked both time and capacity feasibility. To check time feasibility we estimated the round-trip time of a route according to the model specified in subsection 3.3. A route was deemed time feasible if our estimate of round-trip time was no more than eight hours.

The following simple procedure produced routes that are both time and capacity feasible.

1. For each school s use the load e_{st} , its smoothed Class A order volume for each period t , generated earlier.
2. Generate a traveling salesman tour through all 600 schools based on a spacefilling curve heuristic (Bartholdi and Platzman 1988). This tour provides a sequence of schools to be visited.
3. Follow the sequence of schools generated, adding a school to a route if its addition does not violate time feasibility nor truck capacity in *any* time period.

This routing problem is related to the capacitated routing problem (see for example, Bodin and Golden 1981, and Beinstock, Bramel and Simchi-Levi 1993).

6. VALIDATING THE TACTICAL MODEL

We ran the tactical model of the original system with the original five truck routes per day. It was clear that the observed system operation was influenced by additional operational considerations such as poor weather, worker absenteeism, expedited orders, and truck rerouting. However, our model assumes that the inputs such as order entry and truck routes follow the tactical level plan. Thus, we should expect differences between the observed system operation and model output; however, we will demonstrate that the data collected from the actual system operation compares favorably with our model output. The models were all run using LINDO (Schrage 1981) on a Sparc 10 workstation.

Validating Model Input. Even though the underlying demand for Class B items was modeled as steady, the allocation of schools to routes and the uncertain delivery lead times generated by the model caused variance in Class B order volume. The coefficient of variation in the Class B order volume over time generated by the




Figure 5. Histogram of the delivery lead times for the original system generated by the tactical model for the same 30 schools as shown in Figure 1.

model was 0.51. The observed coefficient of variation in Figure 2 was 0.41. Given the difference between the tactical and operational details of the system, we considered this difference to be acceptable.

Validating Model Output. We then compared the lead time generated by the model and that observed in the system. Figure 1 shows the lead time histogram of the 30 sample schools and Figure 5 shows the lead time histogram generated by the model for the same 30 schools. Since the model does not allow operational level route rescheduling, the model-generated lead time histogram of Figure 5 allows only lead times in multiples of two weeks. Even though these histograms differ, the mean of the observed distribution was 20.35 with standard deviation 12.20, while the mean of the model-generated distribution was 21.00 with standard deviation 13.02. Our model-generated lead times closely match historical lead times on the first two moments of the distributions. We conclude that our model is a good representation of the system and therefore can be used to analyze the impact of potential changes.

7. IMPROVED SYSTEM PERFORMANCE: SEPARATING CLASS A AND CLASS B LOGISTICS SYSTEMS

7.1. The Class B Logistics System

The volume of Class B orders over the four-month period was estimated to average just less than one truck load per day. Thus, a separate Class B logistics system with one truck had the potential to deliver Class B items by a single truck to 60 schools daily. This was considered feasible because the CPS system runs interschool mail delivery with a single truck visiting 60 schools per day. Additionally, given the feedback model of Class B orders, our model suggested that consistent delivery of Class B orders would decrease the variance of Class B order volume.

Implementation. In implementing our recommendation to create a separate Class B logistics system, the Director simply added the Class B distribution to the mail delivery system. In this way, schools were able to place an order any day of the week and receive delivery within two days (one day was used to process and pick the order). The feedback model predicts that a decrease in the mean and variance of lead time for Class B items will result in more frequent and smaller orders from the schools. Data collected from September 1993 through April 1994 shows that indeed the number of schools ordering per day had increased by 30% and the average order size had decreased by 25%. The trends show that both of these numbers are improving, but the Director believes it will take some time before schools adjust their behavior. During the same period, the service level (the probability of delivery within two days) to the schools for Class B items improved to over 95%.

7.2. The Class A Logistics System

We then used the tactical model to consider the impact of possible changes in the Class A logistics system. We first ran the tactical model with only Class A items. A new set of routes using four trucks per day was developed by the Traffic Manager. The lead times generated from the model runs indicated a mere 10% improvement in on-time delivery. The mean lead time continued to be close to three weeks.

We next used the tactical model to examine the impact of smoothing the order entry of forecasted engineering Class A items. The model showed on-time delivery improvement to 98.65%. However, truck capacity utilization remained at only 60%.

We next generated new routes using the routing algorithm described in subsection 6.2. In our case, *two* trucks per day was infeasible even with perfect truck packing and balance between routes. This suggested that with a constraint of maintaining the same number of trucks per day, at least three trucks per day are required. Our routing procedure generated a solution with three trucks per day that was both time and capacity feasible.

Implementation. The three routes were implemented and we collected data during the period of February to April 1994. The percentage of days that order volume exceeded truck capacity is shown in Figure 6. At the start of implementation, considerable order volume was backordered at the warehouse which took some time to clear. In addition, the computer-generated routes had to be adjusted for details of the road network, such as major expressways and one-way streets. As Figure 6 shows, the system improved over time until the service level reached 90% in period 5.

8. CONCLUSIONS

The separation of logistics flows provided a means to improve capacity utilization, reduce lead time, and decrease system costs. This was in contrast to the typical notion that pooling items into common logistics systems will reduce overall costs. This separation of logistics flows enabled us to




Figure 6. The Class A logistics system with three daily truck routes improved quickly after implementation.

reduce the three sources of logistics interactions we identified:

1. A *capacity interaction* between Class A and Class B items sharing the same trucking capacity.
2. The *order entry interaction* between the forecasted engineering orders and the order entry system.
3. The *feedback interaction* between the delivery lead times and school orders.

The most critical customers (the school principals) have seen a dramatic decrease in lead times, from over three weeks to just two days. Shortly after implementation of the Class B mail truck delivery, a school survey indicated that 85% of the schools have noticed an improvement in delivery lead times, and 89% claim that they will increase their usage of the central warehouse. About 73% of the schools claim that they can now stock less inventory and depend on the improved service from the warehouse to satisfy demand. The Warehouse Director estimated a decrease in system wide costs of about \$300,000 per year. This was primarily due to a decrease in truck capacity and an increase in order volumes that previously went to third party vendors.

Since completion of this project, the warehouse personnel have been trying a similar approach for food distribution. In addition, since the incoming order streams have been smoothed, a next step to improving warehouse performance is to coordinate delivery by suppliers with shipments by the warehouse. This is expected to decrease warehouse inventory levels and space requirements.

ACKNOWLEDGMENT

We thank Manuel Ortiz, the Director of the Department of Warehousing and Distribution at the Chicago Public Schools, for his enthusiasm and insight into implementing changes to his system. We thank Denise Nash and Claudia Mancini from the Financial Research and Advisory Committee—they were not only our initial contact with the CPS distribution system, but played a vital role of securing data and interfacing with school personnel. We also appreciate the support provided by the Graduate School of

Business at the University of Chicago. Finally, we appreciate the valuable suggestions provided by the Area Editor and the reviewers.

REFERENCES

- BARTHOLDI, J. J., AND L. J. PLATZMAN. 1988. Heuristics Based on Space-filling Curves for Combinatorial Problems in Euclidean Space. *Mgmt. Sci.* **34**, 291–305.
- BEINSTOCK, D., J. BRAMEL AND D. SIMCHI-LEVI. 1993. A Probabilistic Analysis of Tour Partitioning Heuristics for the Capacitated Routing Problem with Unsplit Demands. *Math. O.R.* **18**, 786–802.
- BODIN, L., AND B. J. GOLDEN. 1981. Classification in Vehicle Routing and Scheduling. *Networks* **11**, 97–108.
- BYRNES, J., AND R. SHAPIRO. 1991. Inter-Company Operating Ties. Working Paper **92-058**, Harvard Business School, Boston, Mass.
- FULLER, J. B., J. O'CONNOR AND R. RAWLINSON. 1993. Tailored Logistics: The Next Advantage. *Harvard Bus. Rev.* **May-June**, 87–98.
- HARVARD CASE. 1990. Merloni Elettrodomesticii SpA: The Transit Point Experiment. Harvard Business School, **9-690-003**, Boston, Mass.
- KARMAKAR, U. 1987. Lot Sizes and Sequencing Delays. *Mgmt. Sci.* **33**, 419–423.
- ROSS, S. 1980. *Introduction to Probability Models*. Academic Press, New York.
- SCHRAGE, L. 1981. *Linear Programming Models With LINDO*. The Scientific Press, Palo Alto, Calif.
- SMITH, P. A., AND J. BARRY. 1991. Logistics in Service Industries. Council of Logistics Management. Oak Brook, Illinois.
- ZIPKIN, P. 1986. Model for the Design and Control of Stochastic, Multi-Item Batch Production Systems. *Opns. Res.* **34**, 91–104.