



Consolidation of Customer Orders into Truckloads at a Large Manufacturer

G. G. Brown; D. Ronen

The Journal of the Operational Research Society, Volume 48, Issue 8 (Aug., 1997),
779-785.

Stable URL:

<http://links.jstor.org/sici?sici=0160-5682%28199708%2948%3A8%3C779%3ACOCOIT%3E2.0.CO%3B2-Y>

Your use of the JSTOR archive indicates your acceptance of JSTOR's Terms and Conditions of Use, available at <http://www.jstor.org/about/terms.html>. JSTOR's Terms and Conditions of Use provides, in part, that unless you have obtained prior permission, you may not download an entire issue of a journal or multiple copies of articles, and you may use content in the JSTOR archive only for your personal, non-commercial use.

Each copy of any part of a JSTOR transmission must contain the same copyright notice that appears on the screen or printed page of such transmission.

The Journal of the Operational Research Society is published by Operational Research Society. Please contact the publisher for further permissions regarding the use of this work. Publisher contact information may be obtained at <http://www.jstor.org/journals/ors.html>.

The Journal of the Operational Research Society
©1997 Operational Research Society

JSTOR and the JSTOR logo are trademarks of JSTOR, and are Registered in the U.S. Patent and Trademark Office. For more information on JSTOR contact jstor-info@umich.edu.

©2003 JSTOR



Consolidation of customer orders into truckloads at a large manufacturer

GG Brown¹ and D Ronen²

¹Naval Postgraduate School and ²University of Missouri-St. Louis, USA

We describe the development and operation of an interactive system based on a mathematical optimisation model which is used by a major US manufacturer to consolidate customer orders into truckloads. Dozens of users employ the system daily for planning delivery of orders from manufacturing plants to customers by truckload carriers, saving numerous hours of the users' time and reducing transportation costs.

Keywords: logistics; practice of OR; optimisation

Introduction

In this paper we describe the development and implementation of a computerised system for consolidation of customer orders into longhaul truckloads at a major US manufacturer. The company involved views this new application as a tool providing strategic advantage over its competition, and therefore wishes not to be identified. We have also changed some of the details in the description of the operation and in the data and the results in order to reduce the likelihood of identification. The first section describes the operational environment, and the second section discusses the system design. The third section which describes the system development is followed by a description of the system's operation.

The operation

A major US manufacturer produces several high volume product lines. Each product line is manufactured in a separate set of plants, and for each line there are about half a dozen manufacturing plants. Each product line is also shipped separately. Some plants are fairly close to each other (within less than 100 miles) while others are scattered around the US. Due to the limited shelf life of the products, limited storage space at the plants, and varying, somewhat uncontrollable yield of finished products from raw materials, the availability of finished products at each plant is limited and changes daily. Therefore, each plant may be shipping nationwide, and a customer's order may be split among two or more plants based on product availability and

the customer's proximity to the plants. The company ships hundreds of truckloads of finished products every week to its customers using long-haul contract carriers. The carriers are paid by the miles and number of stops on their assigned routes. Some of these truckloads are for a single delivery destination, whereas others are delivered at multiple locations to different customers. The focus of this work is on those multiple stop routes.

For each product line the US is divided into sales regions. Dozens of sales representatives are selling the products to customers, mainly over the phone. Each sales representative has his own set of geographically clustered customers, and is in charge of full customer service to them. That responsibility includes timely delivery of their orders. Therefore, the sales representative are the ones who consolidate the orders of their customers into truckloads. The product composition and quantities of a customer's order are often changed by the customer before the order is shipped, and, when necessary, the sales representative may negotiate with the customer such changes (within limits) to facilitate consolidation of orders into truckloads. Orders consolidation is viewed by the sales representatives and their supervisors as a chore to keep costs in line. Orders consolidation is done daily for the orders to be shipped on the following day, with visibility of orders several additional days forward for consolidation opportunities.

Some additional features characterize this operation. Order sizes range from 100 lbs. to 45 000 lbs. To assure customer service, consolidated truckloads are limited to at most five delivery stops (the further down the chain the stop is, the less reliable its delivery time). A truck route may span three or even four days. Very often multiple orders exist for delivery from the same plant to the same customer location. A truck may be loaded in two adjacent plants, but the orders loaded at the second plant, which will be at the

tail of the truck, must be delivered first. Due to handling requirements certain orders must be loaded at the nose of the truck and, therefore, delivered at the end of the route.

Before the implementation of the system which we describe here, the order data were presented to the sales rep on a terminal of a mainframe computer, but the consolidation of the orders into truckloads was a manual process using printed maps. The shipment buildup could take several hours per day, cutting significantly into the selling time of the sales rep. The management decided to automate the bulk of the order consolidation process with a primary goal of saving sales rep time and a secondary goal of saving some transportation costs. It was clear that the order consolidation process could not be fully automated because some orders were hard to consolidate, but it was hoped that the bulk of the work could be automated.

This freight consolidation problem falls in the domain of vehicle routing and scheduling problems, a topic which has attracted a significant amount of attention in the literature.¹ However, there are a very large variety of vehicle routing and scheduling problems, and we could not find any work addressing this specific problem. The distinguishing features of this problem are: (a) a truck may be loaded in more than one source for delivery to several locations, (b) the trucks are paid only for a one-way trip to the last delivery location (where they disappear from the system), and (c) flexibility has to be built into the system to accommodate operational dynamics. We elected to use an Elastic Set Partitioning (ESP) model to solve this problem due to its ability to accommodate a wide variety of operational requirements, and the availability of a very efficient solver for this model.

System design

The order consolidation system was designed to assist the sales reps, not to replace them. The sales rep receives a list of recommended consolidations and must approve them, one by one, before they are dispatched. The sales rep still has the option to consolidate the orders manually. Due to frequent changes in the orders, flexibility had to be built into the system to allow the user to experiment with it under different conditions. This was achieved by allowing the user to specify a variety of run parameters and experiment with their values. The order consolidation process is outlined in Figure 1.

For every run initiated by a user three data files are extracted from the corporate database and submitted to the order consolidation system. These are the global file, the plants file, and the orders file (see Table 1). The global and plants files consist of data that seldom change and are under the control of the system administrator. The orders file changes every run and provides data concerning the specific orders to be consolidated and the run parameters as selected by the user.

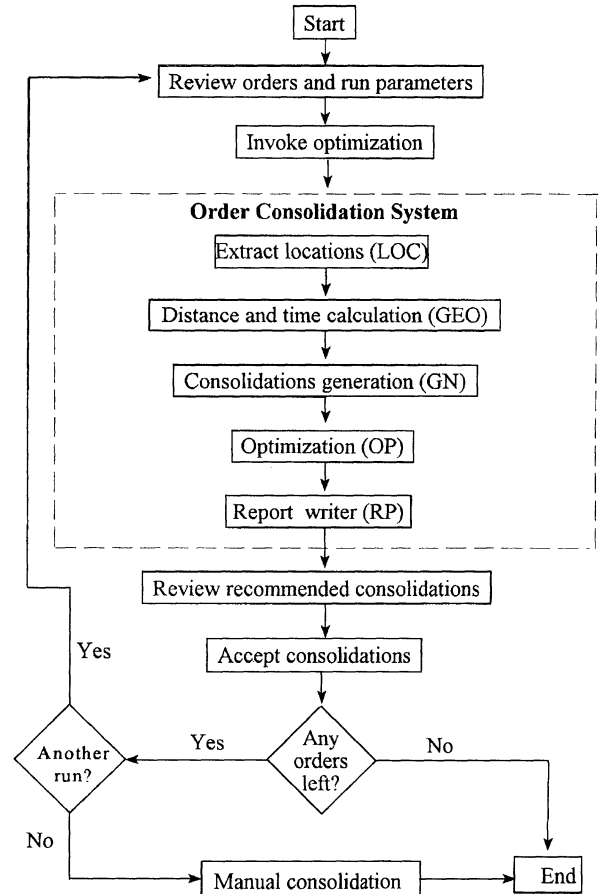


Figure 1 Order consolidation process.

The consolidation system consists of a sequence of five processing steps: extracting locations (shipment origins and destinations) (LOC), constructing a driving time and distance table (GEO), generating potential consolidated loads (GN), optimisation (OP), and report writing (RP). The first step (LOC) extracts from the plants file and the orders file a list of all the locations (plants and order delivery points) which participate in the specific run, and feeds this list into the second step. The second step (GEO) reads the list of locations participating in the run and a file with about 40 000 road segments in the USA, each segment having a distance and driving speed. Using a very efficient shortest route algorithm² a point-to-point distance and driving time table is constructed for the locations participating in the specific run. The third step (GN) involves a data reader and a problem generator. The data reader reads all the three input data files and checks the data for consistency and accuracy. Minor data errors are resolved with default values and warnings, but major ones result in discarding the order involved, or even premature termination of the run. The problem generator creates all legal consolidated loads within the parameters specified in the data. Every combination of orders is considered for a consolidated load, and if the combination is found to be within the specified parameters

Table 1 Input data files*Global file*

Default driving speed
 Driving hours per day
 Default stop time
 Maximal number of delivery stops on a route
 Maximal distance among delivery stops
 Minimal weight for a second plant loading
 Maximal out-of-route miles for each sales region

Plants file

For every plant:
 Plant ID
 Plant latitude
 Plant longitude
 Truck departure time
 Plant operating hours
 Plant type
 Maximal load weight
 Minimal load weight

Orders file

Run parameters:
 Product line
 Source plant(s)
 Destination region
 Can order shipping date be changed? (yes, no)
 Minimal load weight (override)
 Maximal load weight (override)

For every order
 Order ID
 Source plant
 Order weight
 Customer ID
 Delivery location state
 Delivery location latitude
 Delivery location longitude
 Planned shipping date
 Requested delivery date
 Delivery time window
 Order stop time (override)
 Type of delivery date
 Order delivery group

concerning minimal and maximal weight, shipping date, maximal number of delivery stops, and maximal distance among stops, that combination is sequenced. All possible sequences are considered beginning with the shortest one. For each combination of orders the first legal sequence found (considering product loading rules, delivery dates and time windows, and out-of-route miles) is retained as a potential consolidated load. The legality of a delivery sequence is determined by a detailed simulation of every activity on the route, taking into account limitations on driving time. If a legal sequence is not found, that combination of orders is dropped. Thus, all legal potential consolidated loads are generated and passed to the next step along with their route length. The optimisation step (OP) solves an Elastic Set Partitioning problem (ESP, see Appendix) which selects the set of consolidated loads with the minimal total

mileage where any order is not consolidated more than once. The ESP model is solved using our proprietary XS solver³ which is fine-tuned to solve this type of problem. The last step (RP) writes a report which translates the solution into the terms of the user and an output data file. The recommended consolidated loads are presented to the user for action.

System development

The development of the orders consolidation system was a part of a much wider effort of overhauling the corporate information systems. The order consolidation system was selected to spearhead that effort with the expectation that early success will move the whole transition forward. A systems development guidance team was established consisting of representatives of the users, the information systems group, and management. That team met numerous times for many hours to specify the requirements from the order consolidation system, and monitor its development and implementation. The team reviewed the order consolidation process, its interaction with other business systems (especially inventory management and transportation) and was actively involved in designing the order consolidation system and the revised user interface screens. It also decided which parameters of the order consolidation system will be left to the user's control. We were asked to develop the orders consolidation system, based on our earlier experience with similar problems.⁴ A prototype system was developed, and testing with operational data began. These tests resulted in several major changes and enhancements to the prototype which were necessary to increase the credibility and acceptance of the system by the users.

Originally, distances between locations were calculated using great circle distances inflated by a factor for road circuitry. These distances were not accurate enough for short legs. They were replaced by a full road network where distances are calculated by a shortest route algorithm.² The concept of out-of-route miles was introduced in order to control the shape of the consolidated routes. Out-of-route miles are the excess miles of the truck route above the direct distance from the first loading location to the last delivery location. These excess miles are caused by the additional loading or delivery stops on the route. Many carriers will not accept routes with too many out-of-route miles. We reflect the out-of-route miles as a percent of the direct distance, and allow that percent to vary according to the destination region. The rationale for this approach is that consolidated loads going to farther destinations should have lower out-of-route percentage to control the shape of the route and, thus, the transit time. However, for short routes the out-of-route percentage is not controlled.

Initially, only 'good' potential consolidated loads were generated, using a 'sweep' algorithm.⁴ But it turned out that

the complex consolidation rules prevented the generation of some consolidations which were actually used by the sales rep. To establish confidence in the system it was essential to generate all consolidations that would have been considered by the sales reps, and, therefore, we modified the approach to generate all legal consolidations.

At first all orders from the same plant to the same delivery location were automatically assigned to the same truck (if they did not exceed the truck capacity), an approach that seemed to be efficient and improve customer service. But it turned out that this rule is often violated by the sales reps because it may interfere with consolidation of orders to other customers. It is also not necessarily the cheapest way to consolidate the orders.⁵ Therefore, this rule was relaxed, but a feature was added to allow the user to assign such orders to the same delivery group so they will be shipped together. Inconsistent lags between order shipping and delivery dates required allowance of flexibility in the shipping date (reserved inventory may be held a day or two till it is shipped). In addition, rules concerning meeting delivery dates were clarified and delivery dates were classified into several types, based on the requirements of the customers. Also, a few states had strict rules concerning mixing intra- and interstate orders loaded in that state on the same truck, and we had to accommodate them.

Although the system was originally developed for the heavily loaded corporate mainframe computer, because of the slow response times, the system was moved to a dedicated RS6000 computer with data download from the mainframe and results uploading back to the mainframe.

The mainframe input and output user interfaces of the old system were modified by the information systems department to accommodate the additional data requirements for the order consolidation program. However, the mainframe user interface screens are character-based, and the development of a graphical user interface is practically infeasible in this environment. Corporate databases and data interfaces were also changed accordingly by that department. All this was done in parallel to the development of the consolidation system.

An operational system was installed during the third quarter of 1994 and rolled out to the users (the sales reps). Several months later two dozen sales reps were using the system regularly. However, these represent less than half of the potential users.

Adoption of the consolidation system by the users was hampered due to several reasons. The team which guided the development of the system consisted of users which were located close to the corporate headquarters (where the development effort took place), but not of ones from farther

Table 2 Orders data

<i>Order no.</i>	<i>Source ID</i>	<i>Order weight (lbs)</i>	<i>Customer ID</i>	<i>Delivery location latitude</i>	<i>Delivery location longitude</i>	<i>Planned shipping date</i>	<i>Requested delivery date</i>	<i>Delivery time window</i>	<i>Delivery state</i>	<i>Delivery group</i>
1	A	17010	F75	32.78	96.81	19960808	19960810	06-16	TX	
2	A	21797	A39	32.74	97.11	19960808	19960811	06-16	TX	
3	A	3995	S68	38.77	90.37	19960808	19960810	06-11	MO	
4	A	26869	K88	30.31	95.46	19960809	19960812	06-14	TX	
5	A	19443	S98	41.58	93.71	19960809	19960812	07-14	IA	
6	A	2267	F58	29.40	98.51	19960809	19960812	06-11	TX	
7	A	8102	C65	41.00	92.37	19960809	19960812	07-12	IA	
8	A	7239	K95	32.93	97.25	19960809	19960812	06-11	TX	1
9	A	10156	K95	32.93	97.25	19960809	19960812	06-11	TX	1
10	A	4489	F83	29.76	95.36	19960809	19960814	07-14	TX	
11	A	5087	K95	32.93	97.25	19960809	19960812	06-11	TX	1
12	A	3200	K89	32.91	96.64	19960809	19960812	06-16	TX	
13	A	20785	M64	48.23	101.30	19960809	19960812	06-14	ND	
14	A	12542	P33	41.73	93.61	19960810	19960812	05-11	IA	
15	A	29977	F59	33.57	101.83	19960810	19960812	05-11	TX	
16	A	2234	F59	33.57	101.83	19960810	19960812	05-11	TX	
17	A	1310	T18	29.40	98.51	19960810	19960813	06-11	TX	
18	B	8823	T42	42.79	96.17	19960810	19960812	06-15	IA	
19	B	13490	H21	42.01	91.64	19960810	19960812	01-23	IA	
20	B	7887	C61	32.74	97.11	19960810	19960812	06-10	TX	
21	B	8338	C66	43.65	94.46	19960810	19960812	07-16	MN	
22	B	3408	F57	32.91	96.64	19960810	19960812	01-10	TX	
23	B	33223	M65	46.88	96.79	19960810	19960812	06-14	ND	
24	B	5538	T18	29.40	98.51	19960810	19960813	06-11	TX	
25	B	20550	L63	44.88	93.14	19960810	19960812	05-15	MN	
26	B	7032	S64	43.54	96.73	19960810	19960812	06-15	SD	
27	B	3803	S68	38.77	90.37	19960808	19960810	06-11	MO	
28	B	4428	C83	38.61	90.21	19960810	19960812	08-15	MO	

locations. Also, communications and user support are less effective over longer distance. Thus, the adoption rate is higher close to the headquarters. Another problem was lack of interaction with the inventory system (when a user changes the shipping date of an order he must verify product availability for the revised shipping date). Recently, a connection with the inventory system has been established and additional users came on board.

The business environment in which the manufacturer operates is constantly changing, and this necessitates occa-

sional minor modifications in the system, which are usually requested by the users.

System operation

The order consolidation system is installed on a RS6000 computer which is connected to a large IBM mainframe. The user interface is on the mainframe. For each run the three input data files are extracted from the corporate databases on the mainframe and downloaded to the

Table 3 Proposed consolidated loads

Load no.	Delivery stop no.	Location	Day	Hour	Distance (miles)	Drive time	Stop time	Order no.	Order weight (lbs)	Customer ID	Delivery time window	Shipping date	Requested delivery date	Proposed delivery date
1	0	A	1	1.0										
	1	TX	2	6.0	750.2	14.5	2.0	1	17010	F75	06-16	19960808	19960810	19960810
	2	TX	3	6.0	19.8	0.6	2.0	2	21797	A39	06-16	19960808	19960811	19960811
	3	TX	4	6.0	287.8	5.4	2.0	6	2267	F58	06-11	19960809	19960812	19960812
Load totals:					1057.8				41074		Out-of-Route: 4.2%			
2	0	A	1	1.0										
	1	IA	1	5.0	210.7	4.0	2.0	14	12452	P33	05-11	19960810	19960812	19960812
	2	TX	2	10.3	949.5	17.4	2.0	4	26869	K88	06-14	19960809	19960812	19960813
	3	TX	2	13.2	43.1	0.8	2.0	10	4489	F83	07-14	19960809	19960814	19960813
Load totals:					1203.3				43810		Out-of-Route: 21.5%			
3	0	B	1	14.0										
	1	MN	2	7.0	112.1	2.6	2.0	21	8338	C66	07-16	19960810	19960812	19960812
	2	ND	3	6.0	292.5	6.4	2.0	23	33223	M65	06-14	19960810	19960812	19960813
Load totals:					404.6				41561		Out-of-Route: 7.0%			
4	0	A	1	1.0										
	0	B	1	2.9	76.0	1.9	4.0							
	1	IA	1	8.4	56.1	1.5	2.0	18	8823	T42	06-15	19960810	19960812	19960812
	2	SD	1	12.2	84.6	1.9	2.0	26	7032	S64	06-15	19960810	19960812	19960812
	3	ND	2	6.0	496.1	10.2	2.0	13	20785	M64	06-14	19960809	19960812	19960813
Load totals:					712.7				36640		Out-of-Route: 22.2%			
5	0	B	1	14.0										
	0	A	1	15.9	76.0	1.9	4.0							
	1	IA	2	7.0	193.1	3.6	2.0	5	19443	S98	07-14	19960809	19960812	19960812
	2	IA	2	11.6	98.4	2.6	2.0	7	8102	C65	07-12	19960809	19960812	19960812
	3	IA	3	14.3	109.8	2.7	2.0	19	13490	H21	01-23	19960810	19960812	19960813
Load totals:					477.2				41035		Out-of-Route: 111.6%			
6	0	B	1	14.0										
	0	A	1	15.9	76.0	1.9	4.0							
	1	TX	3	6.0	738.1	14.5	2.0	8	7239	K95	06-11	19960809	19960812	19960812
	1	TX	3		0.0	0.0	0.0	11	5087	K95	06-11	19960809	19960812	19960812
	1	TX	3		0.0	0.0	0.0	9	10156	K95	06-11	19960809	19960812	19960812
	2	TX	3	9.3	55.2	1.3	2.0	12	3200	K89	06-16	19960809	19960812	19960812
	3	TX	4	6.0	34.7	0.9	2.0	20	7887	C61	06-10	19960810	19960812	19960813
	4	TX	4	8.9	34.7	0.9	2.0	22	3408	F57	01-10	19960810	19960812	19960813
Load totals:					938.7				36977		Out-of-Route: 20.3%			
7	0	B	1	14.0										
	0	A	1	15.9	76.0	1.9	4.0							
	1	TX	3	5.0	923.4	17.8	2.0	15	29977	F59	05-11	19960810	19960812	19960812
	1	TX	3		0.0	0.0	0.0	16	2234	F59	05-11	19960810	19960812	19960812
	2	TX	4	6.0	390.0	9.0	2.0	17	1310	T18	06-11	19960810	19960813	19960813
Load totals:					1389.3				5538	T18	06-11	19960810	19960813	19960813
Load totals:									39059		Out-of-Route: 34.0%			

RS6000. After the consolidation system concludes the run an output data file is uploaded back to the mainframe and the results are communicated to the user.

The user starts by reviewing the orders and filling a run request screen, specifying orders selection criteria and run parameters under his control: order shipping dates (range), can order shipping date be changed? (yes, no), source plant(s), destination region, delivery dates (range), minimal weight of a consolidation and maximal weight of a consolidation. Several minutes after launching a run (usually less than five) the user receives the results. After the user reviews the recommended consolidations, he may approve all (or some) of them or he may consolidate orders manually. Orders which are not consolidated by the program are dealt with by the user (sales rep) and may require contact with the customer. A run may be launched by the user at any time. Normally, the user launches a run 2–3 times, and that usually results in consolidation of over 90% of the orders. Because there are many users and a single processing machine, submitted runs may have to wait in a queue to be processed, but this has not been a serious problem.

Of the five processing steps, the time consuming ones are the generator (GN) and the optimiser (OP). The generator time depends on the number of orders submitted for the run and on the tightness of the run parameters. The number of orders submitted ranges from less than a dozen to well over a hundred, resulting in from several dozens to many thousands of potential consolidated loads. The run time of the generator is almost always below a minute. The run time of the optimiser depends on the number of potential consolidated loads generated, and is also usually less than a minute. Although all combinations of orders have to be considered in the generation of potential consolidated loads, experience has shown that due to the tight rules and parameters, most of the combinations are rejected upfront. Only a small fraction (less than 10%) are more seriously evaluated, and even the bulk of these are discarded as infeasible. Loose run parameters (especially a large difference between the maximal and minimal weight of a consolidated load) result in generation of a large number of potential consolidated loads. On the other hand, if the run parameters are too tight, for example, allowing a too narrow weight range, a significant share of the orders may not be consolidated. The users have learned this by experimenting with the system parameters.

An example of a (relatively) small set of operational orders is provided in Table 2. This set of orders comes from two plants (sources) and is destined to half a dozen states in the middle of the US. This dataset was run while permitting changes in the shipping dates, and a specified consolidation weight range of 35–45 000 lbs. The recommended consolidations are presented in Table 3. Of the 28 orders in this example, 24 are consolidated into 7 loads, 3 relatively small orders (nos. 3, 27 and 28) could not be consolidated

(no feasible consolidations within the given parameters existed), and one order (no. 25) was not consolidated (although it had 4 feasible consolidations). The first two loads originate at source A, the third one at source B, the fourth load is loaded first in A then in B, and the last three loads are loaded first in B and then in A. One can see that the customers' delivery time windows and the limit on driving time (12 h a day) extend the duration of the routes.

A careful comparison of product shipping costs which was performed by the manufacturer for one of the product lines involved, indicated a 5% cost savings, which translate into 1.7m US dollars per year for that product line.

Conclusion

We have developed and implemented a flexible computerised system for consolidating customer orders into truckloads while minimising truck miles and meeting all customer service requirements. Dozens of sales representatives use the system daily, and it cuts their order consolidation time from hours to minutes, freeing time for selling the product. In addition the system facilitates meeting customer service goals and saves transportation dollars.

In spite of the large variety of operational environments in orders consolidation, we have demonstrated here, and in our earlier work, that in many instances mathematical optimisation models can be implemented and integrated in daily operations. The fast progress of computing and telecommunication technologies further facilitate such applications.

Appendix

The formulation of the ESP model is as follows:

Indices:

- $i = 1, \dots, n$ trucks
- $t = 1, \dots, r$ truck types
- j —an alternate consolidation
- $j \in J(t)$ consolidations requiring truck type t
- $k = 1, \dots, m$ orders to be consolidated
- $j \in J(k)$ consolidations which include order k .

Data:

- C_j —Miles of entire consolidation j (a function of the orders in that consolidation)
- N_t —Number of trucks of type t
- d_t, \bar{d}_t —lower and upper constraint violation penalties for truck type t
- s_k, \bar{s}_k —lower and upper constraint violation penalties for order k .

Decision variables:

- $y_j = 1$ if consolidation j is selected; 0 otherwise
- $\delta_t, \bar{\delta}_t$ —elastic (constraint violation) variables for truck type t

$\underline{\sigma}_k, \overline{\sigma}_k$ —elastic (constraint violation) variables for order k .

Model formulation (ESP):

$$\text{Min } \sum_j C_j y_j + \sum_{t=1}^r (\underline{d}_t \underline{\delta}_t + \overline{d}_t \overline{\delta}_t) + \sum_{k=1}^m (\underline{s}_k \underline{\sigma}_k + \overline{s}_k \overline{\sigma}_k) \quad (1)$$

subject to

$$\sum_{j \in J(t)} y_j + \underline{\delta}_t - \overline{\delta}_t = N_t; \text{ for each truck type } t \quad (2)$$

$$\sum_{j \in J(k)} y_j + \underline{\sigma}_k - \overline{\sigma}_k = 1; \text{ for each order } k \quad (3)$$

$$y_j \in \{0, 1\}; \text{ for each consolidation } j \quad (4)$$

$$\underline{\delta}_t, \overline{\delta}_t \in \{0, 1\}; \text{ for each truck type } t \quad (5)$$

$$\underline{\sigma}_k, \overline{\sigma}_k \in \{0, 1\}; \text{ for each order } k. \quad (6)$$

Each plant is assigned a different truck type because consolidation parameters may be plant-specific. The number of trucks available at each plant is practically unlimited (namely, larger than necessary), therefore we set it equal to the number of orders originating at the plant. Constraint (2) seeks one consolidation for each truck of type t , where a lower violation represents a total idleness of such a truck. Because we are dealing with carrier trucks for which no commitment has been made, the lower violation penalty is set to zero. The upper violation also does not come into play because we provide an excessive number of trucks. Constraint (3) seeks to consolidate all orders. A lower violation represents an unconsolidated order (and is set to the round trip mileage from the source plant) and an upper violation is assigned a high disruption penalty (in order to prevent it from occurring).

A typical (large) problem may have over a hundred orders originating in several plants, which translate into over a hundred constraints and tens of thousands of columns (binary variables). The ESP problem is solved to within 0.1% of optimality (the solution to the relaxed LP is a lower bound for the objective function value).

The elastic feature of this set partitioning model is essential for this application. Without the possibility of constraints violation (at a cost) a feasible solution would usually not be available (due to the tight constraints specified by the users). Even if one ignores the operational limitations (for example, out-of-route miles, shipping and delivery dates) it is very often impossible to consolidate a set of orders into truckloads weighing 40–45 000 lbs.

References

- 1 Assad AA (1988). Modeling and implementation issues in vehicle routing. In: Golden BL and Assad AA (eds). *Vehicle Routing: Methods and Studies*. North-Holland, Amsterdam, pp 7–45.
- 2 Brown GG and Halwachs TE (1993). *Making Transportation Models Believable and Visible*. Presented at ORSA/TIMS Conference, Phoenix, AZ, USA. (Available from author)
- 3 INSIGHT (1990). *The X-System: A Large-Scale Linear, Integer, and Nonlinear Optimization System*. Insight Inc. Alexandria, VA.
- 4 Bausch DO, Brown GG and Ronen D (1995). Consolidating and dispatching of truck shipments of mobil heavy petroleum products. *Interfaces* **25** (2): 1–17.
- 5 Dror M and Trudeau P (1990). Split delivery routing. *Naval Res Logist* **37**: 383–402.

*Received August 1996;
accepted April 1997 after one revision*