Collaborative Logistics: The Shipper Collaboration Problem

Özlem Ergun* Gültekin Kuyzu* Martin Savelsbergh*

*Industrial and Systems Engineering and The Logistics Institute, Georgia Institute of Technology
Atlanta
GA 30332, USA
{oergun,gkuyzu,mwps}@isye.gatech.edu

1 Introduction

Shippers and carriers are continuously facing pressures to operate more efficiently. Traditionally shippers and carriers have focused their attention on controlling and reducing their own costs to increase profitability, i.e., improve those business processes that the organization controls independently. More recently, shippers and carriers have focused their attention on controlling and reducing system wide costs and sharing these cost savings to increase every one's profit. A system wide focus, e.g., a collaborative focus, opens up cost saving opportunities that are impossible to achieve with an internal company focus. A good example is asset repositioning. To execute shipments from different shippers a carrier often has to reposition its assets. Shippers have no insight in how the interaction between the various shipments affects a carrier's asset repositioning costs. However, shippers are implicitly charged for these repositioning costs. No single participant in the logistics system controls asset repositioning costs, so only through collaborative logistics initiatives can these costs be controlled and reduced.

A good example that reveals the benefits of collaboration involves scheduling truckload movements of multiple shippers. Consider a tour set up between two of the participants in the Nistevo collaborative logistics network. Through collaboration, they have created and are using a dedicated 2,500-mile continuous move tour that visits many Midwestern and Eastern U.S. cities. The seven legs of this tour encompass various distribution centers, production facilities, and retail outlets for both companies. The tour minimizes asset repositioning costs and gives the carrier's drivers a repeatable schedule. This tour resulted in a 19% savings for both shippers over their former company-centric model. At the same time, the carrier is experiencing higher margins through better asset utilization and lower driver turnover through more regular driver schedules. Collaborative logistics creates a true win/win scenario for both shippers and carriers.

However, identifying routes minimizing asset repositioning costs in a collaborative logistics network is not easy. Especially, when the number of participants in the network, and thus the number of truckload movements, grows, the number of potential routes to examine becomes prohibitively large. In that case, optimization technology is needed to assist analysts in identifying tours with little or no asset repositioning.

The core optimization problem, called the lane covering problem (LCP), is stated as follows: given a set of lanes, find a set of tours covering all lanes such that the total cost of the tours is minimized. More formally, given a directed Euclidean graph $D = (V, A)$ with node set V, arc set A, and lane set $L \subseteq A$, find a set of simple directed cycles covering the lanes in L of minimum total length. Let l_{ij} denote the length of arc (i, j) and let x_{ij} be an integer variable indicating how often arc (i, j) is traversed. Then the solution to the following minimum cost circulation problem can be decomposed into a set of simple directed cycles covering all lanes with minimum total length:

$$z_L = \min \sum_{(i,j)\in A} l_{ij}x_{ij}$$

$$s.t. \quad \sum_{j\in N} x_{ij} - \sum_{j\in N} x_{ji} = 0 \quad \forall i \in N \tag{1}$$

$$x_{ij} \geq 1 \quad \forall (i, j) \in L \tag{2}$$

$$x_{ij} \geq 0 \quad \forall (i, j) \in A \setminus L. \tag{3}$$

In our problem definition, we have implicitly assumed that a single truckload has to be moved across each lane. It is simple to handle shipments consisting of multiple truckloads. If v_{ij} denotes the number of shipments that has to be moved across lane $(i, j) \in L$, then all we have to do is replace $x_{ij} \geq 1 \ \forall (i, j) \in L$ with $x_{ij} \geq v_{ij} \ \forall (i, j) \in L$ in the formulation above.

In earlier research, we have studied variants of the lane covering problem in which the length of a cycle, C, is constrained (either $|\{(i, j) : (i, j) \in C\}| \leq k$ or $\sum_{(i,j)\in C} l_{ij} \leq K$). We have developed heuristics which produces near optimal solutions within seconds for instances with up to 1000 lanes.

In our current research, we extend the initial work in two directions (both motivated by practical complications): we take into consideration the timing of shipments, and we study fair cost allocation (gain sharing) mechanisms.

2 Timing of shipments

Solutions to the lane covering problem as defined above may not be implementable due to various timing issues. The yearly freight volume on a lane translates into a frequency of shipments. In practice, these frequencies can vary from once every two weeks to twice a day. Clearly, if drivers have to wait a long time between moving shipments on two consecutive lanes in a cycle, truck and driver utilization go down and the costs for the carrier go up. To further complicate the situation, drivers are subject to stringent Department of Transportation (DoT) regulations that limit driving hours and duty period hours. Drivers cannot drive more than 11 hours in a duty period, the duty period cannot be longer than 14 hours, and there have to be at least 10 hours of rest between consecutive duty periods (according to the new Hours of Service regulation going into effect in January of 2004). These temporal aspects make the construction of tours with high utilization more complicated, where utilization is defined as the fraction of the duration of the tour spend on actually moving freight.

To analyze the effect of timing issues on shipper collaboration, we define and study the following extension of the lane covering problem (denoted by LCP-T). Given a directed Euclidean graph $D = (V, A)$ with node set V, arc set A, travel times t_{vw} for arcs $(v, w) \in A$, lane set $L \subseteq A$, and time windows $[e_v, l_v]$ on the dispatch time at the origin for lanes $(v, w) \in L$, find a set of simple directed cycles covering the lanes in L of minimum total duration. Minimizing total duration ensures high resource utilization, where utilization is defined as the ratio of the sum of travel times on the lanes in a cycle and the duration of the cycle (the difference between end time and start time).

To introduce an even higher degree of accuracy, we also study a variant in which we consider the impact of drivers. We enforce that each tour has to be such that it can be executed by a driver, i.e., without violating DoT regulations governing driving hours and duty period hours, assuming that the driver is available at the start of the tour at the start of the driver's duty period.

Initially, we assume that the planning period covers one week, i.e., all time windows on the dispatch time for the lanes fall within the planning period and that all lanes repeat weekly. Consequently, during tour construction, we will allow tours that start on a Friday and end on a Tuesday.

We formulate LCP-T as a set covering problem and generate a feasible solution using a fast greedy heuristic. Let $\mathcal{C}_K$ represent the set of all time feasible directed cycles in D. Let u_C denote the utilization of cycle C, l_C be 1 if lane l is on the cycle C, and x_C be a 0-1 variable indicating whether cycle C is selected or not. Then LCP-T can be formulated as follows

$$\min \sum_{C \in \mathcal{C}_K} u_C x_C$$

$$\sum_{C \in \mathcal{C}_K} l_C x_C \geq 1 \quad \forall l \in L$$

$$x_C \in \{0, 1\} \quad \forall C \in \mathcal{C}_K.$$

It is important to observe that the complexities introduced by the timing of shipments are completely isolated in the construction of time feasible directed cycles.

A key difference between generating cycles to cover lanes with windows on the dispatch time and generating cycles to cover lanes without windows on the dispatch time is that it is now possible that a cycle $((i_1, i_2), (i_2, i_3), ..., (i_{n-1}, i_n))$ is infeasible, but that the cycle $((i_2, i_3), (i_3, i_4), ..., (i_{n-1}, i_n), (i_1, i_2))$ is feasible. Furthermore, since windows on the dispatch time are specified within a planning period and are assumed to repeat, we have to consider cycles that "wrap around" the end of the planning period during cycle generation. Efficient implementations have been developed that specifically handle these complications. Extensive computational results will be reported based on real-life instances.

3 Fair cost allocations

To be successful in practice, it is important that the rules governing the collaboration are specified and agreed upon by all members in advance. A key aspect, if not the key aspect, of such an agreement is how the gains achieved by collaborating are distributed among the members. Distributing the gains is equivalent to allocating the costs. Therefore, we have investigated *fair* cost allocations. More precisely, given the total cost for covering all the lanes how should the total cost be allocated among the shippers?

First, we consider cost allocation mechanisms that are *stable*, i.e., no subset of the members has an incentive to opt-out and collaborate only among themselves, and *budget balanced*, i.e., the total cost of covering the lanes is allocated among the members. In cooperative game theory terminology such cost allocation mechanisms are said to be in the *core*.

There is a close relation between the core and linear programming duality. For certain NP-hard problems, e.g., facility location, it has been shown that the core is empty and, furthermore, that the integrality gap is equal to the gap between the recoverable and the total cost.

We show that the unconstrained version of the lane covering problem has a cost allocation that is in the core. Moreover, such a cost allocation can be determined in polynomial time by solving the dual of the minimum cost circulation formulation of the lane covering problem given above.

Without loss of generality, we assume that every shipper has only one shipment. Let α_{ij} be the payment that is allocated for covering lane $(i, j) \in L$. If the cost allocation determined by α_{ij} is in the core then the following conditions must hold:

$$z_L = \sum_{(i,j)\in L} \alpha_{ij} \tag{4}$$

and

$$\sum_{(i,j)\in S} \alpha_{ij} \leq z_S \quad \forall S \subseteq L, \tag{5}$$

where z_S is the total cost of optimally covering the lanes in S.

Now consider the dual of the LP in (1-4), where u_{ij} are the dual variables for constraints (2) and (3), and y_i are the dual variables for the constraints (1):

$$d_L = \max \sum_{(i,j)\in L} u_{ij}$$

$$s.t. \quad u_{ij} + y_i - y_j = c_{ij} \quad \forall (i,j) \in A \tag{6}$$

$$u_{ij} \geq 0 \quad \forall (i,j) \in A. \tag{7}$$

Theorem 1 *The cost allocation defined by* $\alpha_{ij} = u_{ij} \quad \forall (i,j) \in L$ *is in the core.*

Before proving Theorem 1, we state two properties related to the given cost allocation scheme:

Property 1 Complementary slackness conditions imply that for any $(i, j) \in A \setminus L$ if $x_{ij} > 0$ then $u_{ij} = 0$, that is given any cycle the cost of covering a cycle is allocated to the non-deadhead arcs.

Property 2 Complementary slackness conditions imply that for any $(i, j) \in L$ if $x_{ij} > 1$ then $u_{ij} = 0$, that is if a lane arc is covered more than once then its cost allocation will always be 0.

Proof of Theorem 1. Condition (4) is satisfied by the LP duality theorem. Next, we show that condition (5) is also satisfied.

Given any set $S \in L$, let A_S be the arcs in the optimal cover of the lanes in S when only the lanes in S are present, including the replicas of the arcs that are used more than once. A_S can be decomposed into single cycles. Then summing the dual constraints (6) over all the arcs in A_S we obtain

$$\sum_{(i,j)\in A_S} (u_{ij} + y_i - y_j) = \sum_{(i,j)\in A_S} u_{ij} = \sum_{(i,j)\in A_S} c_{ij} \quad \forall S \subseteq L. \tag{8}$$

Note that y_i's cancel out because A_S is decomposed into simple cycles. For the problem restricted to the lanes in S the total cost incurred is equal to:

$$z_S = \sum_{(i,j)\in A_S} c_{ij} \quad \forall S \subseteq L. \tag{9}$$

By (8) and (9) we get

$$\sum_{(i,j)\in A_S} u_{ij} = z_S \quad \forall S \subseteq L.$$

The total cost allocated for covering the lanes in set S in the original solution is equal to

$$\sum_{(i,j)\in S} \alpha_{ij} = \sum_{(i,j)\in S} u_{ij} \quad \forall S \subseteq L,$$

since by Properties 1 and 2, $\alpha_{ij} = 0$ for all deadhead arcs and for all lane arcs that are used more than once in the original solution. We can re-write $\sum_{(i,j)\in A_S} u_{ij}$ as

$$\sum_{(i,j)\in A_S} u_{ij} = \sum_{(i,j)\in S} \alpha_{ij} + \sum_{(i,j)\in A_S \setminus S} u_{ij} = z_S \quad \forall S \subseteq L. \tag{10}$$

By (7) $u_{ij} \geq 0$ for all $(i, j) \in A$, hence (10) implies

$$\sum_{(i,j)\in S} \alpha_{ij} \leq z_S \quad \forall S \subseteq L. \qquad \square$$

It is also easy to show that the core for the shipper collaboration problem is always non-empty, since both the primal and the dual of the lane covering problem are always feasible.

Unfortunately, finding a cost allocation mechanism in the core does not directly imply that this is a reasonable model for the application at hand. For example, in the cost allocation model described above it is possible to have members with 0 payments if their lane is traversed more times than required. Furthermore, there are simple examples where no cost allocations in the core without this property exist. For example, consider 3 shippers, 2 of them sending a shipment from A to B and one from B to A. Let the cost of traversing the arc (A, B) be 1. Then any cost allocation assigning $\epsilon > 0$ for covering lane B to A cannot be in the core.

A *population monotone* (or cross monotone) cost allocation mechanism has the property that when the number of members in the collaboration increases the individual cost shares do not increase. Cross monotonicity implies fairness. In the shipper collaboration problem it is important to device population monotone allocation schemes. It is easy to show that the allocation scheme above is not population monotone. Furthermore no allocation schemes in the core can be population monotone.

We are investigating (and will report on) the existence of population monotone allocation mechanisms and the effect of using membership fees and utility functions.