Strategies for Coordinated Drayage Movements

Christopher Neuman and Karen Smilowitz*

May 9, 2002

Abstract

The movement of loaded and empty equipment (trailers and containers) between rail yards and shippers/consignees is a costly part of intermodal shipping. These drayage costs are exacerbated by the lack of coordination among parties, including shippers, railroads, trucking companies, and intermodal marketing companies. This paper reports initial findings regarding the design of strategies to coordinate drayage movements, with a focus on the Chicago freight interchange, a region unique in its size and complexity of operations. A simple heuristic, used to obtain feasible solutions and provide insight into specialized algorithms, is presented. This research is part of a larger study on Chicago intermodal freight operations.

1 Introduction

It has been estimated that sixty percent of all container movements in the United States enter or move through the Chicago region, see Rawling (1997). These movements consist mainly of drayage operations (i.e., loaded and empty trips between rail yards and shippers/consignees), see Li *et al.* (2001). Drayage is a costly part of intermodal shipping. Forty percent of a 900 mile movement cost is incurred in the drayage portions, typically less than 50 miles, see Morlok and Spasovic (1994). These drayage costs are exacerbated by the lack of coordination among parties, including shippers, railroads, trucking companies, and intermodal marketing companies. This paper presents initial findings regarding the design of strategies to intelligently coordinate the movement of trailers and containers. This work focuses on the Chicago freight interchange, a region unique in its size and complexity of operations.

Section 1.1 describes the Chicago drayage problem (CDP) in detail. Section 1.2 presents literature related to the CDP. Section 2 introduces the notation necessary to represent the CDP. The solution approach is presented in Section 3, along with initial findings. Conclusions and future work are discussed in Section 4.

1.1 The Chicago drayage problem

Drayage movements both outbound from a shipper to an intermodal yard and inbound from an intermodal yard to a consignee are illustrated in Figure 1 with a simple example.[1] In Figure 1(a), a tractor bobtails (i.e., repositions without a container) to the container yard to pick up a container. The container is then moved to a shipper who has requested a pickup. The driver then has two options: he can either wait at the shipper until the container is loaded ("stay with") or depart for another destination and leave the container at the shipper ("pick and go"). If the pick and go option is chosen, a second tractor later bobtails to the shipper to transport the loaded container to the intermodal yard. The container is then loaded on to a train for long haul transportation.

*Department of Industrial Engineering and Management Sciences, Northwestern University

[1] Drayage operations involve both containers and trailers. For ease of illustration, only containers are considered here, recognizing that trailers can be modeled in a similar manner with the same solution techniques.

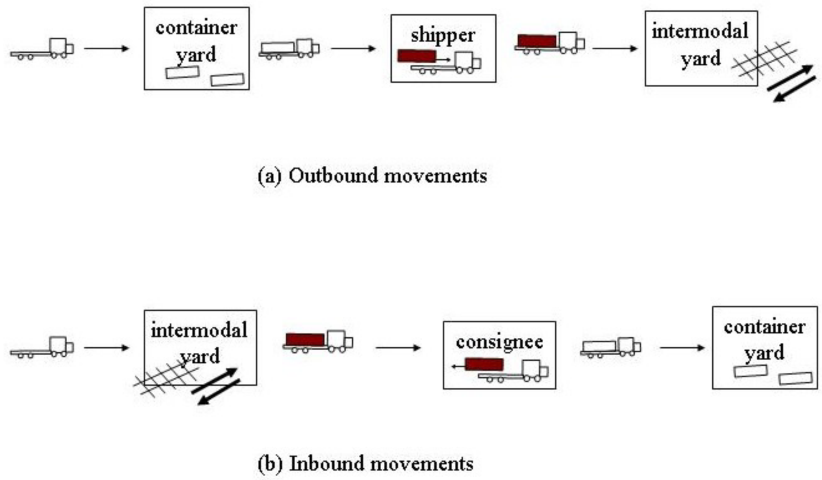

(a) Outbound movements

(b) Inbound movements

Figure 1: Drayage movements

Figure 1(b) depicts the movement of a container to the consignee. A tractor bobtails to intermodal yard to pick up the loaded container and deliver it to the consignee. A driver can then either wait for the container to be unloaded or leave the container at the consignee. Again, if the pick and go option is chosen, a second tractor is required to return the empty container to a container yard. All movements must adhere to specified time windows for pickups and deliveries, as well as hours of operation for container yards.

Alternative routing to reduce bobtailing and empty container movements exist. Rather than returning the container to the yard, the driver could reposition the container to a shipper that has requested an empty container. The extent to which drayage costs can be reduced through coordination can be measured by the reduction in empty movements and bobtailing. When measuring these savings, congestion effects within the network should be taken into account. The system as a whole can operate better if bobtailing and empty container movements occur during off-peak hours. Therefore, the savings measures include incentives to perform these activities during off-peak times.

Currently, the assignment of drivers to movements is controlled by several independent parties, most commonly by intermodal marketing companies. In this project, we will explore the possible savings that can be achieved through coordination of drayage operations among the various parties. Here we consider the extreme case where a single entity controls operations to gain insight into the potential savings.

1.2 Related literature

Since limited work exists on the modeling of drayage operations, more general work in the routing of tractors and trailers is considered as well. Morlok and Spasovic (1994) presents an in-depth look at drayage operations for Conrail. The drayage problem, described in Morlok and Spasovic (1990), is modeled with a network flow formulation. Time windows on pickups and deliveries are addressed by discretizing time over the planning horizon. The network formulation contains a large number of variables and constraints, which limits the application of this model to larger problems. In Ball *et al.* (1983), the trailer distribution for a chemical company is studied. A network flow formulation, similar to the work in Morlok and Spasovic (1990), is compared with a second formulation where origins and destinations are aggregated and a single node represents the trailer trip requested. The problem is then transformed into a vehicle routing problem (VRP) where tractor tours are created to serve requested trailer trips. As a result, standard heuristics for the VRP are applied, including a savings algorithm and a greedy insertion method. Like Morlok and Spasovic (1990), the network flow formulation is limited to smaller problems. In De Meulemeester *et al.* (1997), a similar VRP formulation is applied to the collection and delivery of "skips". The skips, comparable to containers and trailers in the drayage problem, are transported between multiple locations. Again, typical VRP solution methods are applied.

Bodin *et al.* (2000) presents a method to coordinate the distribution of tractors and trailers in trash disposal operations, termed the "Rollon-Rolloff Vehicle Routing Problem" (RRVRP). The objective of the RRVRP is to find a set of routes to minimize miles traveled by tractors to serve all trips. A trip is defined as the empty and loaded movements required to serve a customer's request. Route lengths are constrained by the length of a work day. The RRVRP is modeled as an asymmetric VRP. Solution methods are developed to exploit the characteristics of various trip types. For example, trips that begin and end at the depot require no additional deadheading and can thus be "packed" into tours of other trip types with a modified bin-packing algorithm. The specialized solution methods, including a decomposition algorithm with modified bin-packing, a partial enumeration method, and a trip insertion/improvement method, are compared with a less complex Clarke-Wright savings method. The partial enumeration heuristic performed consistently well. The performance of the savings method varied with trip type mix, which is to be expected since the other heuristics are based on the characteristics of the trip types.

The Chicago drayage problem poses several new issues. The most important factor is simply the size of the problem, in terms of the number of shippers, consignees, and yards, and the total movements per day. The time window constraints on pickup and delivery lead to additional complications. Further, models must incorporate peak time congestion. One must consider truck behavior at shippers and consignees (i.e, stay with and pick and go movements). This behavior can be either input to the model or a variable specified by the model. The combination of these factors makes the network formulation in Morlok and Spasovic (1990) difficult to implement. The heuristics developed in Bodin *et al.* (2000) appear promising; however, the increased complexity of the Chicago drayage problem may pose problems with route generation algorithms. Section 3 presents the first steps in developing new solution methods for the Chicago drayage problem.

2 Network description

In this section, the notation to describe the drayage network is presented. As in Bodin *et al.* (2000), customer requests and the subsequent tractor movements are modeled as trips. The set $\mathcal{T}$ contains all trip types, including both the set of loaded and empty trips, $\mathcal{T}_l$ and $\mathcal{T}_e$, respectively. The specific trip types are presented in Figures 2 and 3 with their required bobtailing. The trips in the figures assume that no joining of trips to form vehicle routes has occurred. Therefore, as the figures show, all trips begin and end at a single point, denoted "TE" for "tractor entry". This point represents a virtual location at which all tractors enter the system.

Set of loaded trips

1. Full container from shipper to rail yard; return empty to container yard [**stay with**]
2. Full container from shipper to rail yard [**drop and pick**]
3. Full container from rail yard to consignee; return empty to container yard [**stay with**]
4. Full container from rail yard to consignee [**drop and pick**]

Set of empty trips [**all drop and pick**]

1. Empty container from container yard to shipper
2. Empty container from consignee to container yard
3. Empty container from consignee to shipper

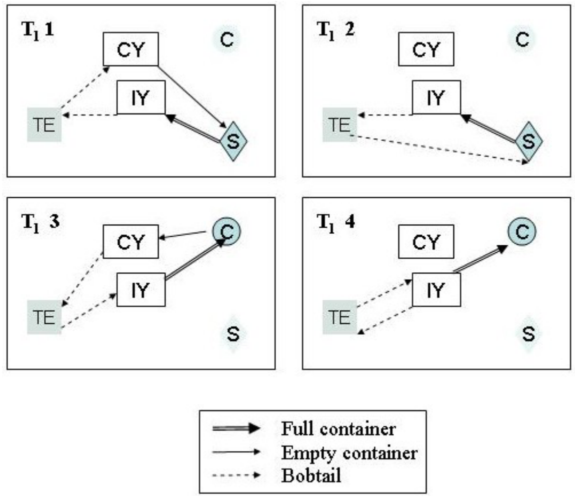

Figure 2: Trip types in drayage operations: loaded trips

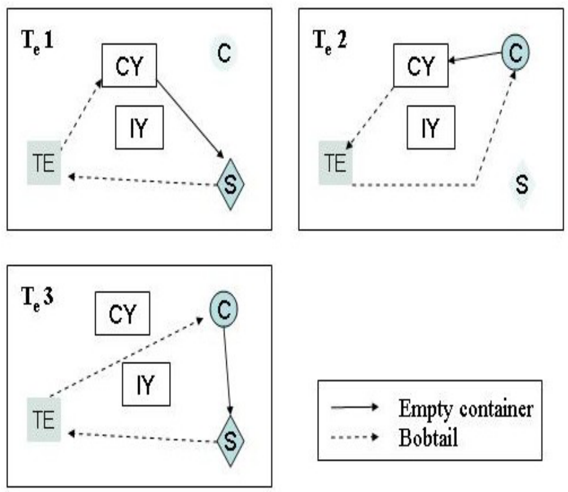

Figure 3: Trip types in drayage operations: empty trips

The Chicago drayage system to be modeled contains multiple intermodal and container yards, as well as multiple shippers and consignees. Individual trips are defined by origin, destination and time window. Perfect information and time-dependent demand are assumed.

3 Solution approach

The objective of the Chicago drayage problem is to assign trips, as defined in the previous section, to vehicle routes at the minimum cost. As discussed in Section 1.2, new solution methods are required to solve the Chicago drayage problem. We believe that the difficult task of generating vehicle routes should be addressed first. For problems of realistic size, enumerating all feasible vehicle routes that can serve trips is impossible, due primarily to the size of the empty trip set. Ultimately, this research will explore algorithms to generate a reasonable subset of routes. As a first step, a modified savings method is presented. This modified savings method provides insight into route generation algorithms. The solutions provided by the savings method should suggest ways to limit the set of vehicle routes to be generated. In addition, a simple heuristic, like the savings method, is a good way to benchmark proposed solution methods in terms of complexity, solution time and quality of solutions.

Modified savings method

Step 0: Define trip set, $\mathcal{T} = \mathcal{T}_l \cup \mathcal{T}_e$

Step 1: Create savings matrix

- Compute $\mathcal{T} x \mathcal{T}$ savings matrix
  1. savings, $s_{ij} = f(i,j)\ \forall i,j \in \mathcal{T}$
  2. function of deadhead and bobtailing trips created and eliminated (with discount for off-peak travel)
- Rank elements in descending order

Step 2: Join first feasible set of trips

Step 3: Update savings matrix

- Remove infeasible empty movements
- Recalculate savings
  1. $s_{mk} = f(m,k)$ for merged trip, $m = i, j, .., \in \mathcal{T}$; $k \in \mathcal{T}$

Step 4: Repeat 2 and 3 until no further trips can be joined

4 Results and discussion

The modified savings method is applied to test problems to evaluate the performance of the modified savings algorithm and the potential savings from drayage coordination. Results are compared with a base case where there are no drop and pick movements and all containers must return to the yard.

4.1 Implementation issues

Computational results from a series of small test cases are presented. The test cases were designed to span an array of driver behaviors. Therefore the percentage of total trips predetermined to be stay with operations varies from 100% to zero. In cases where no trips are predetermined as stay with, driver behavior becomes a decision variable.

As shown in Figure 4, the Chicago region is idealized to be an area 25 miles square. Eight customers (four shippers requiring pickup and four consignees requiring delivery) are distributed uniformly in the area, as are three rail terminals and two container yards. Customers were randomly allocated to a rail terminal and a container yard with uniform probability. Distances were calculated using the Euclidean distance metric, and converted to travel times using an average speed of 30 mph, which incorporates slack for unexpected stochastic travel delays.

A driver serving a stay-with customer must wait one hour for loading or unloading a container, while a driver serving a drop and pick customer takes fifteen minutes to load or unload a container. It takes fifteen minutes to load or unload a container at a container yard, and thirty minutes to load or unload a full container at a rail terminal. The savings algorithm was tested on the same data, but the mix of stay with and drop and pick customers varies from 0% (all drop and pick) to 100% (all stay with). It was assumed that all information about customer type, assignment, and location were known in advance. The maximum tour length is ten hours, equivalent to a single day shift.

The savings algorithm starts by assuming that a separate tractor, originating and terminating at a depot or "virtual terminal", serves each customer. In this example, the virtual terminal is located

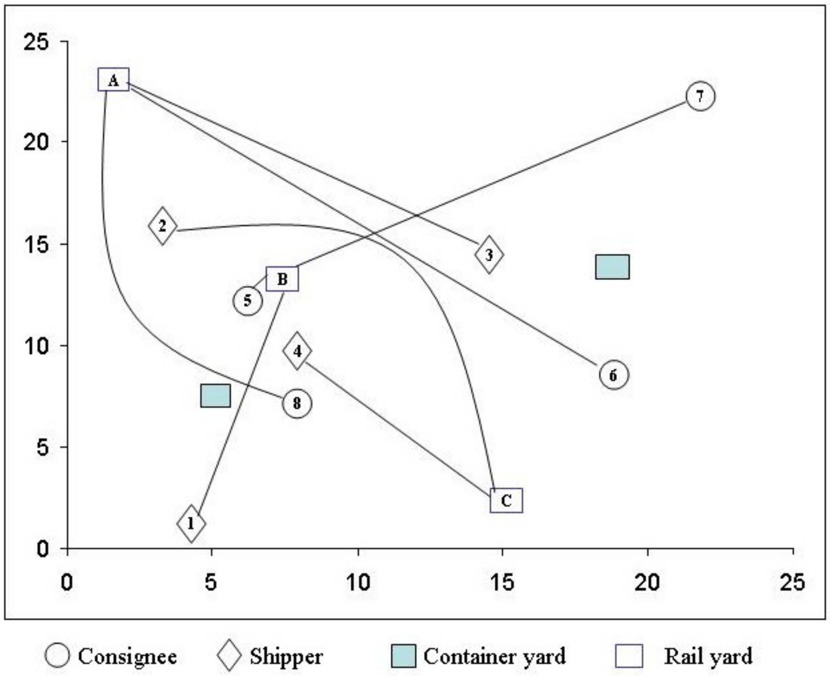

Figure 4: Trips in test case

at the origin, i.e. in the southwest corner of the region. Savings were calculated for every possible sequencing of two customers, and were asynchronous (i.e., the savings achieved by sequencing trip 1 after trip 2 was possibly different from sequencing trip 2 after trip 1). In this example, time windows were not incorporated.

4.2 Computational results

Results for the four runs are presented in Table 1, and compared with results when the eight trips were served by eight separate tours. The required fleet size was measured by the number of tours needed to serve all trips. Two tours were created for three of the four runs, while three were required when all trips were stay with. The tours were not necessarily balanced, with the 100% run having one tour of only five hours (the maximum was ten) and the 0% run having a tour of four hours, half that of the other tour for the same run. The number of customers served per tour increases as the percent of stay with customers decreases, because there are fewer trips required to or from the container yards.

The tours were constructed manually, and the process of checking for valid savings in the savings matrix led to some interesting insights about the tours' construction. The first major insight is that opposite tour types attract. Sequencing pickup trips after delivery trips (and vice versa) resulted in the most savings, as opposed to scheduling pickups after pickups or deliveries after deliveries. This is intuitive, since a delivery ends at a container yard and a pickup starts at a container yard (in the case of stay with). Conversely, the pickup ends at a rail yard and a delivery starts there. As a result, if two trips share a container yard (rail yard) that is a long distance from the virtual terminal, the savings will be significant, since the cost of a long deadhead back to the terminal is saved while the distance between the endpoint of one trip and the starting point of the next is zero. It is important to note that the model did not consider the additional savings achieved by sequencing a stay with pickup after a stay with delivery, which would allow a truck to bypass the container yard entirely.

As more potential merges became invalid, there were merges of existing trips in such a way

	Total tour time	Fleet size required	Average trip time
Separate trips	19.0	8	2.4
Coordinated trips			
SW level			
All trips	23.8	3	7.9
63%	18.6	2	9.3
25%	15.3	2	7.6
No trips	12.9	2	6.4

Table 1: Savings from drayage coordination

that pickups were sequenced after pickups, and likewise for deliveries. Converting these tours into a series of "opposite" sequencings could result in significant savings. This is a feature that more sophisticated heuristics will be able to exploit.

Examining the method by which potential merges became invalid provided additional insight into route structure. For example, if the highest-ranked savings was to sequence delivery customer 6 after pickup customer 3 (as was the case in all runs), it was a relatively mechanical process to run through the matrix and determine which merges were no longer valid. No sequencing of trip 3 before any other trip was valid; the sequencing of trip 3 after trip 6 was invalid; and no sequencing of anything before 6 was valid. More generally, sequencing a trip y after an existing tour $\{x_1, x_2, ..., x_n\}$ would render all sequences of x_n with any other trip invalid, would render invalid any sequencing of trip x_i after y, and would allow no other x_i to be sequenced before trip y. This information will be of great use in coding updates to the saving matrix, as well as when time windows are incorporated.

5 Conclusions and future work

Although the Chicago drayage problem involves many complications described in Section 1.1, we were able to gain key insights for future solution methods from the implementation of the savings algorithm with a simple test case. In particular, the issues raised about route generation will be beneficial in future steps. Using these insights, new solution heuristics beyond the modified savings algorithm can be developed. For example, a set covering approach will be explored in the future. The savings algorithm can be used to benchmark performance of new methods.

Acknowledgments

Partial funding was provided by the National Science Foundation, grant EHR-0125752.

References

Ball, M., Golden, B., Assad, A., and Bodin, L. (1983). Planning for truck fleet size in the presence of common-carrier options. *Decision Sciences*, **14**(1), 103–120.

Bodin, L., Mingozzi, A., Baldacci, R., and Ball, M. (2000). The roll-on, roll-off vehicle routing problem. *Transportation Science*, **34**(3), 271–288.

De Meulemeester, L., Laporte, G., Louveaux, F., and Semet, F. (1997). Optimal sequencing of skip collections. *Journal of Operations Research Society*, **48**, 57–64.

Li, I., Ziliaskopoulos, A., and Waller, S. (2001). Modeling intermodal freight flow movements in the chicago area. *Transportation Research Record*.

Morlok, E. K. and Spasovic, L. N. (1990). An optimization model of tractor-trailer drayage operations. Research report, University of Pennsylvania.

Morlok, E. K. and Spasovic, L. N. (1994). Approaches for improving drayage in rail-truck intermodal service. Research report, University of Pennsylvania.

Rawling, G. (1997). Statistical summary and value of the intermodal freight industry to northeast Illinois. Research report, Chicago Area Transportation Study.