Operational Support for Regional LTL Carriers

Alan L. Erera* Martin W.P. Savelsbergh* Prashant Warier*

*School of Industrial and Systems Engineering
Georgia Institute of Technology
Atlanta, GA 30332-0205
`{alerera, mwps, pwarier}@isye.gatech.edu`

1 Introduction

The trucking industry is vitally important to the U.S. economy, providing essential dock-to-dock and dock-to-home freight transportation services. We develop operational planning decision support methodology for *regional less-than-truckload (LTL) carriers*, a segment of the trucking industry focused on fast, reliable transport of less-than-truckload lot sizes.

LTL carriers serve customers that ship quantities ranging from 150 lbs to 10,000 lbs, i.e., less-than-truckload quantities. To be economically viable, LTL carriers must consolidate smaller shipments into nearly-full truckloads. A truckload typically consists of two 28-foot pup trailers. To enable consolidation, carriers maintain cross-dock terminal networks with *end-of-line* terminals serving customers in relatively small geographic areas (e.g., a city) and *break-bulk* terminals serving end-of-lines in relatively large geographic areas (e.g., a state). This hub–and–spoke structure of the so-called *linehaul network* is necessary because there is often not enough freight to support near-truckloads between end-of-line terminals.

Linehaul network operations refers to transportation supporting freight movements from origin end-of-lines to destination end-of-lines, potentially moving through one or more break-bulk terminals. Linehaul networks also typically include *relay* locations with trailer marshaling yards; most relay locations correspond to break-bulks. When the travel time between terminals is too long for service by a single driver trip, new drivers may be substituted at relays. Although individual trailers are not unloaded at relays, they may be recombined with other trailers before being dispatched (known as *pup-matching*) and thus relays play a role in freight routing.

Fundamental to LTL linehaul operations is the *load plan*. For freight moving between each origin end-of-line and destination end-of-line, the load plan specifies a (usually unique) sequence of break-bulk terminals at which that freight will be cross-docked. In the default case, freight is typically handled at two break-bulks: an origin break-bulk and a destination break-bulk. To reduce handling costs, carriers often add *direct loads* to the plan which allow freight to skip handling at either the origin or destination break-bulk. While an increase in direct loads leads to decreased average freight transportation time, it usually does not lead to decreased line-haul transportation cost, since trailers are typically routed via the same sequence of terminals whether or not they are moving direct. Of course, execution of a load plan requires subsequent trailer and driver assignment decisions.

There are two primary types of LTL carriers in the U.S.: *national* carriers which operate nation-wide networks, and *regional* carriers which serve more limited market areas. National and regional LTL carriers compete with one another for business. Due to a larger market, national carriers handle much more freight and thus have more opportunity for consolidation, which generally leads to lower handling and linehaul transportation cost per unit. Since regional carriers cannot compete with nationals on price, they attempt to provide higher service levels (i.e., shorter transit times). National carriers typically operate static load plans which are infrequently adjusted during operations. Due to less-stable demand patterns and more restrictive service requirements, regionals must be more flexible and alter their load plans frequently.

To ensure high service levels and to limit costs, regionals often employ *meet-and-turns*. A meet-and-turn provides the same functionality as a relay, i.e., driver exchange and pup-matching, but at dynamically-selected locations such as parking lots and/or highway rest areas. A carefully selected meet-and-turn can reduce transportation costs by eliminating some of the inherent circuity of the linehaul network. Furthermore, effective use of meet-and-turns can reduce driver costs by reducing layover costs, i.e., costs incurred when a driver does not return to his domicile at the end of his duty. For example, drivers traveling in opposite directions on a lane may meet somewhere near the travel time center, exchange loads, and return to their starting points.

The need for cost-effective dynamic planning is a unique challenge for regional LTL carriers and has been the focus of our research. We have developed a set of models that, on a daily basis, determine how to route freight through the system, including the selection of appropriate meet-and-turns, how to build and dispatch truckloads, and how to assign drivers to loads, so as to minimize costs. Initial computational results have been promising.

2 Operational support for regional LTL carriers

The dynamic nature of operations at regional LTL carriers makes the availability of effective decision support tools for daily planning activities more important. Presently, regional LTL carriers make changes to the load plan, but in an ad hoc manner and at the discretion of the terminal managers at end-of-lines and break-bulks. This may result in faster service, but, since decisions are based on local information only, may come at the expense of using more trailers and drivers than are actually needed to move the freight. Centralized daily planning tools can address this shortcoming, while additionally mitigating the potential problems generated by seasonal variations in freight volumes under static load plans.

In this research, we develop a set of 24-hour planning models for regional LTL linehaul network operations. The three planning models determine in sequence:

1. *Freight routing*: the selection of daily meet-and-turn locations and the assignment of origin-destination freight flows to paths through the linehaul network;
2. *Trailer assignment and dispatch timing*: the assignment of flows to trailer-loads given a

limited set of flow paths and the determination of dispatch times for trailer pairs;

3. *Driver assignment*: the assignment of drivers to truckloads for dispatch.

Although a single, integrated model might be preferable since it could account for all system costs simultaneously, computational limitations on practically-sized instances prohibit this approach.

The planning models are designed to be executed each afternoon at the completion of pickup-and-delivery operations. At this time, the carrier has a set of new freight shipments at end-of-line terminals, and a set of incomplete multiple-day shipments at break-bulks. We assume that each shipment has a unique overnight destination. In cases when the final end-of-line destination for a shipment cannot be reached overnight, we assume that an external load plan is used to determine its overnight destination. Shipments at each terminal with a common destination are aggregated into a splittable flow volume.

The models require specification of the underlying linehaul network, including fixed break-bulk and end-of-line terminal locations as well as potential meet-and-turn locations. Travel legs connecting these locations are specified with distance-based cost and travel time. Note that only legs that can be completed in a single driver shift are included. The driver assignment model also presumes complete driver information, with current location, availability time, remaining working hours, and domicile for each driver.

3 The freight routing model

The first solution phase focuses on the assignment of freight flow volumes to paths through the linehaul network, minimizing the sum of transportation and handling costs. Transportation costs accrue per truckload-mile, while handling costs accrue both for freight cross-docking and for trailer pup-matching. We assume unlimited handling capacity at all locations. Further, to simplify the initial routing model we relax the problem by assuming that: (1) any freight arriving at a break-bulk can be consolidated with any other freight, regardless of arrival time; (2) any trailers can be pup-matched with any other at a terminal, relay, or meet-and-turn regardless of arrival time; and (3) an unlimited number of drivers and trailers are available at each location. These three relaxations will be removed in the subsequent models.

The freight routing model is solved as a multi-commodity network flow problem. A commodity is defined for each freight origin-destination pair. The primary decision variables include (1) the freight volume of commodity ℓ flowing over transportation leg (i, j); and (2) the total integer number of truckloads flowing over leg (i, j). Volume decision variables are specified as semi-continuous variables of the form $x_{ij}^{\ell} \in \{0\} \cup [a_{\ell}, \infty)$ to prevent paths with too little volume. In addition to standard freight conservation constraints, we include simple consolidation constraints to model the number of truckloads required to transport the assigned flow on each leg, with consolidation allowed between freight commodities.

Distinguishing features of our freight routing model is that it incorporates dynamic selection of meet-and-turn locations and explicitly determines pup-matching decisions. To model pup-matching, we introduce additional network nodes at each terminal and potential meet-and-turn location. At each such pup-matching node, we include standard freight conservation constraints, and introduce additional constraints that both enforce *trailer* conservation and prohibit shipments from moving from one trailer to another. To do so, we make one simplifying assumption which is not limiting in practice: all freight flow with destination d moving through pup-matching location m has a unique next stop k. Note that k may be d or a break-bulk serving d. Trailer variables account for the number of trailers inbound to location m with next stop k.

The freight routing model has been solved using data supplied by a major U.S. regional LTL carrier. The primary instance considered includes 13 break-bulk terminals, 70 end-of-lines, and 98 potential meet-and-turn locations. Freight flows are given for 787 origin-destination pairs. To limit the set of potential freight routing options, potential intermediate break-bulk and pup-matching locations for each origin-destination pair are preselected within an elliptical area; all practical locations are included. The resulting instance is an integer program with 9458 rows and 19,496 columns.

To enable solution in reasonable computational time, we implement a specialized branching scheme. Branching priority is given first to the truckload flow variables and then to freight volume flow variables. Initial computational experiments, using XPRESS-Optimizer 2003 on a 1.6GHz Mobile Pentium 4, are promising; we are able to find integer solutions within 10% of optimality in less than 10 minutes. Ongoing computational work is directed at improving solution speed and quality.

4 The trailer assignment and dispatch timing model

After solving the freight routing problem, we have information on how the freight between two end-of-lines should be routed through the linehaul network, i.e., the paths and associated volumes to use. However, the freight routing model assumes that freight arriving at a break-bulk can be consolidated with any other freight, regardless of when the freight arrives. In reality this is not possible, since trailers arrive at the break-bulk terminals at different times throughout the operating period. Additionally, pup-matching decisions at non-terminal meet-and-turn locations may be operationally infeasible since extended wait times may violate driver hour constraints or create late loads. Consequently, more trailers (and hence drivers) may be necessary than predicted by the routing model, and alternate freight paths may in some cases be desirable.

The trailer assignment and dispatch timing model is a more detailed version of the freight routing model. It minimizes the same objective, the total transportation cost required to move freight through the linehaul network. Given the freight flows assigned to paths identified by the routing model, this model builds trailers and two-trailer truckloads for dispatch, explicitly accounting for freight *ready* and *cut* times and segment travel times. Given a specific path through the linehaul network, the latest dispatch, or cut, time for freight with a specific destination can be easily calculated at each location. The freight earliest dispatch, or ready, time is known at the freight origin. At downstream locations in the path, however, the ready time depends on earlier dispatch times.

At non-terminal meet-and-turn locations, a driver must wait for the driver he is meeting to arrive before completing the pup-matching task and continuing. Since driver hour constraints prohibit excessive waiting times at these locations, we include constraints that limit this wait time. When timing considerations are introduced, some meet-and-turn and relay decisions made by the freight routing model are no longer desirable. We therefore introduce a single alternate path for each such freight flow, and introduce binary indicator variables to model the selection of the preferred path.

The trailer assignment and dispatch timing model is a time-indexed integer programming formulation to assign freight to trailers and truckloads, while ensuring that no ready and cut times are violated. The primary complexity arises from the fact that precedence constraints which model restrictions on dispatch time choices on consecutive legs in an freight flow path. Computational experiments have shown that the increase in cost from the freight routing model is modest, and varies between 4-8%.

5 The driver assignment model

In the last phase, we assign available drivers to truckloads minimizing total driver costs, which include driving costs, non-driving costs and layover costs. Assigning drivers to loads is complicated by the fact that loads have a dispatch time window (implied by the latest ready time of a shipment in the load and the earliest cut time of a shipment in the load) and because there exist precedence relations between loads since some individual shipments move from one load to another at break-bulk locations.

As a first step, we decide for which terminal-terminal legs to use meet-and-turns, in the true meaning of the term, to allow drivers to return to their domicile. We only consider legs for which the travel time is such that it is impossible for drivers to make an out-and-back trip and for which there are appropriately timed opposing loads (loads going in each direction). Let a and b be the terminals, let m be the meet-and-turn, and let tt_{am} and tt_{bm} be the relevant travel times. Furthermore, let r_a, r_b, c_a, c_b, d_a, d_a, t_a, and t_b denote the ready time, the cut time, the driver ready time, and the yet-to-be-determined dispatch time for the loads at terminal a and b. We decide to use a meet-and-turn when the following set of inequalities has a feasible solution:

$$\begin{aligned} max(r_a, d_a) \leq t_a \leq c_a \\ max(r_b, d_b) \leq t_b \leq c_b \\ |(t_a + tt_{am}) - (t_b + tt_{bm})| \leq t' \end{aligned}$$

where t' is parameter which limits the time drivers are allowed to wait for each other at the meet-and-turn location. (Any feasible solution is acceptable, but we prefer the one that minimizes $t_a - max(d_a, r_a) + t_b - max(d_b, r_b)$.) The above discussion extends easily to meet-and-turns where pup-matching operations take place.

We assign drivers to the remaining loads using one of two approaches: (1) a time-discretized integer programming formulation covering the entire planning period, or (2) an iterative greedy optimization-based heuristic. The integer program tends to get large and difficult to solve quickly because of the need to model precedence relations between loads. The iterative greedy heuristic avoids this complication by working with time periods short enough to ensure that two loads with a precedence relation between them are never considered during the same iteration. For each time period an assignment problem involving drivers and loads is solved. We consider drivers that are available during the time period, loads with a cut time in the time period (assignments of drivers to these loads are forced), and loads that are ready (or will become ready) during the time period but with a cut time in subsequent periods (assignments of drivers to these loads are encouraged, but not forced). Only feasible assignments are allowed, e.g., a driver must be able to reach the origin of a load before its cut time, and a driver must be able to reach the origin of a load and be able to complete the loaded move within his remaining drive/duty hours. Critical to the success of the iterative greedy heuristic is defining appropriate costs for each feasible assignment. We include the following in the cost of an assignment: number of resulting non-loaded driving hours, number of resulting non-driving hours, benefit for returning to driver domicile (inverse proportional to driving hours left upon return), benefit for moving a load which does not have its cut time in the current time period, and benefit for moving a load early. After the assignment problem for a time period has been solved, the dispatched loads are removed from the problem and the driving/duty hours left for drivers are updated. Information from dispatched loads is used to update the ready times of subsequent loads.

Extensive computational experiments are conducted to demonstrate the efficacy of the proposed approach for providing operational support to regional LTL carriers.