

A MATHEMATICAL MODEL FOR DRIVER BALANCE IN TRUCKLOAD RELAY NETWORKS

Sarah Root

University of Arkansas

Hector A. Vergara

University of Arkansas

Abstract

Driver retention has been cited as one of the primary motivating factors for the implementation of relay networks for full truckload transportation. The strategic design of such networks considering important operational factors such as limitations on load circuitry and equipment balance has been previously studied in the literature, however driver scheduling has not been explicitly considered in routing decisions. We present a prescriptive modeling approach that uses mathematical programming in conjunction with a decomposition-based algorithm to select feasible duties that consider current hours-of-service regulations and assign them to drivers domiciled at relay points in the network to cover truckload demands during a given planning horizon. Computational results are presented for randomly generated problem instances along with areas for future research.

1 Introduction

Driver retention is one of the most significant challenges in full truckload (TL) transportation [1]. TL carriers commonly use a direct point-to-point (PtP) dispatching system for the movement of freight; that is, loads are picked up at their origin and moved to their destination by a single driver. Under this system, drivers spend a significant amount of time on the road given the long distances that they need to cover and the difficulty finding appropriate backhaul trips. Some estimates put this number at two or three weeks at a time [1]. As a result, drivers perceive a reduction in their quality of life and they tend to quit; typical driver turnover rates for TL carriers exceed 100% annually [2]. For this reason, driver retention has motivated the analysis of alternative dispatching systems for TL transportation. One of the alternatives is to install a network of relay points (RPs) where drivers and trailers can be exchanged. A network of RPs would allow truckloads to continue moving to their final destinations, and the drivers to return home

more frequently as described in [3], [4], [5] and [6]. This would resemble the operations of LTL carriers except that freight would not need to be sorted at RPs [5]. Driver turnover rates for LTL carriers are significantly lower than for TL carriers, approximately 20% annually [2].

The configuration of a relay network (RN) would be similar to a hub-and-spoke network. In this network, each truckload must stop at one or more RPs while being transported from its origin to its destination. For example, Figure 1 shows an illustration of a RN where truckload ij is moved from origin node i to destination node j through a series of three relay points.




Figure 1: Relay Network for Truckload Transportation

The strategic design of these networks considering important operational restrictions such as limitations on load circuitry, equipment balance and driver tour length has been previously studied in [5], [6], and [7]. However, driver scheduling has not been explicitly considered in routing decisions, and the basic assumption made up to this point is that once the RN is designed and the routes for the truckloads are selected, the assignment of truckloads to drivers domiciled at the RPs can be made easily and in a way that allows them to return home frequently.

However, there are several factors that require driver scheduling to be considered explicitly. First, hours-of-service regulations for the trucking industry enforce specific rules that are important when determining how to assign drivers to loads. The recent addition of new rules for driver management makes this an even more challenging problem for the carriers. Second, driver shortage has been consistently listed as one of the top six concerns for the trucking industry during the past seven years [8]. The use of relay networks for TL transportation will only help to alleviate this problem if the assignment of drivers to truckloads is explicitly considered along with other strategic and tactical decisions for the purpose of obtaining appropriate driver tour lengths. This is important to retain drivers, especially since that there is an estimated shortage of 125,000 drivers in 2011 as reported in [9].

The issue of scheduling drivers is relevant to other logistics operations that use network based configurations. There is potential to extend this research to consider the needs of less-than-truckload (LTL) carriers and third party logistics providers, as well as to provide a means to assess the effect on cost and performance of utilizing shared

facilities and transportation equipment as proposed by collaborative logistics efforts like the Physical Internet initiative [10].

In this paper, we present a new modeling approach for selecting *duties* – a series of loads to be moved and rest periods between movements for drivers that originate and terminate at a particular domicile – and the assignment of these duties to drivers based at RPs in a relay network for TL transportation. One of the contributions of our work is that we conduct a preliminary analysis to explore the effect of freight lane volume and time away from domicile allowed for drivers on the characteristics of the solutions obtained. This is an important step in understanding the tradeoffs of operating using a relay network structure as compared to current PtP dispatching. Also, having a model for incorporating these important tactical decisions in TL dispatching planning is another step towards making relay networks applicable in practice, and understanding how relay networks can be incorporated in collaborative logistics efforts.

The remainder of this paper is organized as follows. In Section 2, we review previous research in relay network design and driver scheduling in the trucking industry. In Section 3, we present the technical approach for the selection and assignment of duties for drivers to cover truckload demand. Section 4 describes the computational experiments completed to assess the performance of our proposed approach and analyze the effect of different characteristics of this problem on the solutions obtained. Next, in Section 5, we highlight the major findings of our research and conclude by discussing areas for future work in Section 6.

2 Problem Description

2.1 Literature Review

The scheduling of drivers in the trucking industry depends on the dispatching system used. TL carriers commonly dispatch loads assigning one driver to a single load from origin to destination using PtP dispatching. On the other hand, LTL carriers use a hub-and-spoke system in which drivers are assigned to smaller loads with multiple origins and destinations and use the hubs for sorting or consolidation [1]. Although these two types of operations are intrinsically different, there are some studies in the literature that focus on the design of relay networks for TL transportation. The motivation behind these studies is to improve driver retention by using a configuration that would allow drivers to return home more frequently. Most of the early work in relay network design developed descriptive simulation analyses of hub-and-spoke networks and alternative dispatching methods for TL transportation such as the ones presented in [1], [3], and [4]. These studies explored different strategic and operational aspects of RN design and demonstrated the feasibility of such systems. More recently, prescriptive models have been proposed for this problem. Üster and Maheshwari [5] and Üster and Kewcharoenwong [6] propose a mathematical formulation for the strategic design of a TL relay network and develop heuristic and exact solution methods. However, important

operational constraints such as limitations on load circuitry are relaxed and the modeling approach is intractable for realistically sized problem instances. Vergara and Root [7] propose a composite variable model for the design of relay networks that implicitly incorporates the difficult operational constraints in the definition of the variables. In particular, the variables used in their model represent feasible truckload routes; this modeling approach allows them to solve largely-sized problem instances efficiently. A similar modeling approach can be used for scheduling drivers needed at the RPs of the resulting relay network to account for hour-of-service regulations and the difficult cost structures that exist in the driver scheduling problem.

To this point, the majority of the research on driver scheduling in trucking has considered LTL operations. Erera et al. [11] present a scheme for the dynamic management of drivers for a major LTL carrier that combines greedy search with enumeration of time-feasible driver duties. Their approach is capable of generating driver schedules that satisfy several operational driver constraints efficiently. In Erera et al. [12], the authors assign drivers to home domiciles in an LTL trucking terminal network. They use an iterative scheme to allocate drivers to domiciles and to determine drivers' bids while satisfying hours-of-service regulations and union rules that restrict driver schedules. Finally, Erera et al. [13] provide a computational approach for the creation of operational schedules for the tactical load plans that are used by an LTL carrier. The scheduling approach presented in this research creates dispatches for loaded trucks between terminals with specified time windows, and then covers all dispatches using cyclic schedules for drivers. The authors emphasize the idea that developing detailed operational schedules allows the estimation of operational costs for a given load plan more accurately along with the evaluation of important performance metrics. All of these studies reinforce the idea that driver scheduling is a very challenging optimization problem, due largely to the challenges of incorporating operational restrictions such as hour-of-service regulations and the difficulty estimating the costs needed for the model.

A recent contribution in driver scheduling in the context of TL trucking is the work of Goel and Kok [14]. They consider the traditional PtP dispatching system and provide a model for scheduling drivers according to the sequence of time windows for the loads that are included in a tour. Although they explicitly consider hour-of-service regulations in their model, their work is only applicable for the PtP dispatching system and does not account for the new rules that have been recently announced to go into effect starting in the second half of 2013 [15].

2.2 Problem Statement

Although the problem we propose to solve is motivated by the TL trucking industry, we anticipate that our model can be used by other systems that utilize drivers to move freight through hub-and-spoke networks. As such, the problem is defined as follows, given a relay network and the associated freight movements between relay points, determine the number of drivers required at each domicile RP, and how to assign drivers to loads

without violating operational constraints such as hour-of-service regulations, a carrier-established limitation on the time away from domicile allowed for drivers, and service requirements to ensure on-time delivery of loads. This problem is solved within the context of TL trucking as described in the following section.

3 Technical Approach

Our proposed technical approach decomposes the driver scheduling problem in a series of smaller problems that are solved sequentially. The following subsections describe the characteristics of our proposed approach.

3.1 Decomposition Approach

Our approach assumes that time-sensitive freight in a hub-and-spoke network must be transported within pre-specified time windows. Recall that our motivating case is TL relay networks, but such problems arise in other contexts as well. Our procedure assumes that the loads and their time windows are each given as input, and that driver duties must be devised and assigned to drivers to transport this freight. Figure 2 shows the main building blocks of our technical approach.

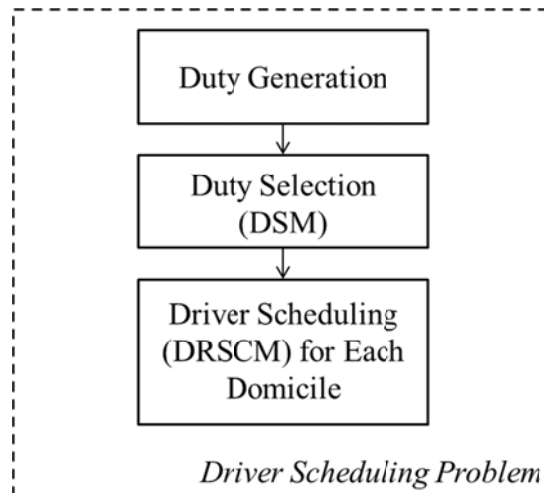


Figure 2: Decomposition Algorithm for Driver Scheduling in Truckload Relay Networks

The driver scheduling problem has three main subproblems: the generation of duties, the selection of duties to be used in the solution, and the assignment of these duties to individual drivers. To generate duties, information about the loads and their time windows is used as input. Recall that duties are a sequence of loads that an individual driver transports; this sequence of loads must begin and terminate at a driver's home domicile. The feasibility of a driver duty depends both on hour-of-service

regulations and a carrier-established limitation on the number of days a driver can be away from his or her from domicile. We refer to the latter limitation as time away from domicile (TAFD), and in this research explore the effect of this important design parameter since it dramatically affects driver retention because excessive TAFD negatively affects a driver's quality of life. We detail the generation of driver duties in Section 3.2.

These duties then become the variables in our modeling approach that selects a set of optimal duties needed to satisfy the demand in the network while minimizing operational costs. This optimization model is called *Duty Selection Model* (DSM), and is described further in Section 3.3.1.

Since duties must begin and end at a driver's home domicile, the selection of duties partitions the duties and their corresponding loads into those that begin and end at each domicile. We therefore can consider how to assign drivers to each of the individual domiciles independently. To do this, we develop a mathematical model that is used to solve the *Driver Scheduling Model for Each Domicile* (DRSCM). This is described in Section 3.3.2. This model determines the minimum number of drivers required at each RP and assigns selected duties that start at this RP to these active drivers during a specified planning horizon. Once the number of drivers required at each RP is obtained, the operational cost of the system can be estimated using a fixed cost for the active drivers and the routing and lodging cost previously obtained from DSM.

3.2 Duty Generation

The generation of duties is the first step to solving the driver scheduling problem. The duties that we generate will become the variables that are subsequently used in the models introduced in Section 3.3. In this research, we use an enumeration-based procedure for the generation of our variables. For this purpose, we define *templates* – predefined duty patterns that can be assigned to drivers domiciled at a RP in the relay network. When enumerating these templates, the combination of loaded movements and empty/bobtail movements in a single duty need to satisfy service requirements for the loads that are being transported in addition to the hours-of-service regulations to generate a feasible duty. We only generate a duty if it can satisfy the time window requirements for earliest dispatch and latest arrival of the loads included in it.

We implemented an algorithm to check the feasibility of each of these templates for each combination of loads in the network. If the feasibility requirements are not satisfied for a given template, then those variables are not included in our model for duty selection (DSM). By considering hour-of-service regulations, the limitation on time away from domicile and the satisfaction of service requirements implicitly within the definition of the variables, we do not need to include them as constraints in our mathematical formulation of DSM.

We use two primary types of templates to generate duties: *out-and-back* and *triangle* templates. Out-and-back templates account for movements between two relay points,

while triangle templates consider visiting three relay points. Figure 3 shows an out-and-back template for a duty that includes two loaded movements: one from domicile RP i to a RP at node j and the other back from j to i . Depending on the distance between nodes i and j and the time windows of the loads, this template type is able to generate two different types of duties. A *short* duty is generated if the travel time between i and j allows a driver to drop-off a load at node j and pick-up another load at this location that will arrive at node i before exceeding a limit on the number of hours that can be driven in one workday, ω . Alternatively, a *long* duty of this template is generated if the load from j to i is dispatched after the driver rests for a sufficient time (τ) as required by the hours-of-service regulations at node j . These long out-and-back duties keep drivers away from their home domicile for two days. Note that we can create an alternative template by replacing one loaded movement with an empty/bobtail movement to balance drivers in the relay network in case there are no backhaul loads. To create duties with empty/bobtail movements we need to check the feasibility of the duty by computing the travel time required between the two relay points using the Euclidean distance and an average speed for the trucks.

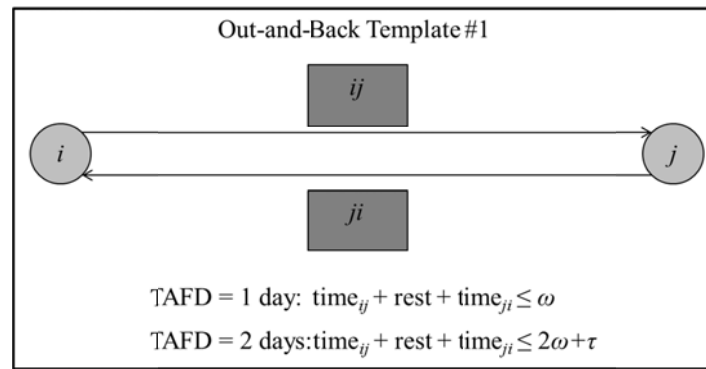


Figure 3: Example Template for Duty Generation

The number of templates needed to enumerate duties depends on a limitation on the TAFD. This is a design parameter that must be established by the carrier. In the present work, we consider templates that are able to generate duties that keep drivers away from their domiciles for up to 3 days. In our computational results in Section 4, we discuss the effect of different values of TAFD on the performance of our approach and the characteristics of the solutions obtained. Figure 4 shows the template types that we used in this research.

In the algorithm that we implemented to generate duties, the cost associated with each feasible duty is computed based on the miles driven and the number of rests away from domicile (i.e., lodging expenses or sleeper berth use compensation). Having a cost estimate for each complete duty is another advantage of our proposed approach. Other types of formulations would require building a duty through a mathematical model that

makes several decisions and has difficulty capturing the non-linear cost structures that are present in this problem.

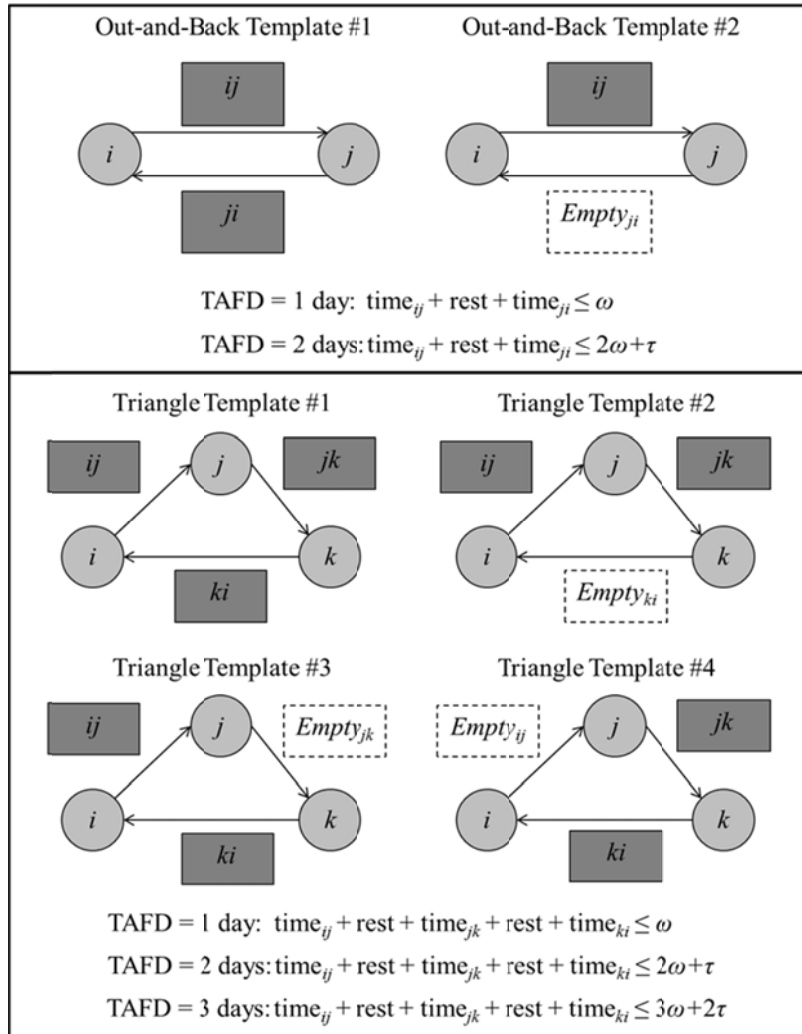


Figure 4: Set of Templates for Duty Generation

3.3 Model Formulation

The proposed decomposition approach and the mathematical formulation of DSM and DRSCM assumes that the duties have been generated as described in Section 3.2, and uses the following notation:

Sets

N = set of nodes k

T = set of truckloads t

L = set of loads l , $|L| \geq |T|$

D = set of duties d

J = set of time periods j

I_k = set of drivers i at node k

$D(l)$ = set of all duties d that contain a load l , $D(l) \subseteq D$

$D_k(j)$ = set of all duties d that start at node k and occur during time period j , $D_k(j) \subseteq D$

$Q_k(d)$ = set of duties that start at node k and are incompatible with duty d , $Q_k(d) \subseteq D$

D' = set of duties d that represent extended rests for drivers, $D' \subseteq D$

Parameters

Φ = length of planning horizon (in number of time periods)

c_d = cost of duty d , $\forall d \in D$

c_i = cost of activating driver i , $\forall i \in I_k$

Variables

$x_{id} = \begin{cases} 1 & \text{if driver } i \text{ is assigned to duty } d \\ 0 & \text{otherwise} \end{cases}$

$y_i = \begin{cases} 1 & \text{if driver } i \text{ is active} \\ 0 & \text{otherwise} \end{cases}$

$z_d = \begin{cases} 1 & \text{if duty } d \text{ is selected} \\ 0 & \text{otherwise} \end{cases}$

3.3.1 Duty Selection Model (DSM)

In this section, we present the mathematical formulation for the problem of selecting driver duties to cover truckload demands while minimizing operational costs in the network (i.e., Step 2 of our proposed algorithmic approach). This approach assumes we have the set of duties generated as described in Section 3.2.

$$\min \sum_{d \in D} c_d z_d \quad (1)$$

subject to

$$\sum_{d \in D(l)} z_d = 1 \quad \forall l \in L \quad (2)$$

$$z_d \in \{0,1\} \quad \forall d \in D \quad (3)$$

This is a set partitioning formulation where the objective function (1) minimizes the total cost, including transportation costs and lodging for drivers who spend a required rest period away from their domiciles. Constraint (2) enforces the satisfaction of load

demands across the network. Finally, Constraint (3) requires all variables in the model to be binary. Recall that since duties have been generated to ensure feasibility given that hours-of-service regulations, carrier requirements, and service requirements for the loads are implicitly considered during the generation of the variables as described in Section 3.3, we do not need to include these constraints in the DSM explicitly.

Recall that each duty begins and ends at a domicile, and contains one or more loads that must be transported. Therefore, the selection of duties essentially makes a unique assignment of loads to drivers who are based at a specific domicile. Since each domicile has a collection of loads that must be moved by drivers at that location, the problem can be decomposed to consider each domicile individually without loss of optimality. Our solution methodology therefore uses the solution obtained from DSM as input for the scheduling of drivers at each domicile in the network using the mathematical formulation for DRSCM presented in the following section.

3.3.2 Driver Scheduling Model for Each Domicile (DRSCM)

Assuming D_k^* represents the set of optimal duties selected in DSM that start at domicile k , the mathematical formulation for scheduling drivers based at domicile k follows.

$$\min \sum_{i \in I_k} y_i \quad (4)$$

subject to

$$\sum_{i \in I_k} x_{id} = 1 \quad \forall d \in D_k^* \quad (5)$$

$$\sum_{d \in D_k^*(j)} x_{id} = 1 \quad \forall i \in I_k, j \in J \quad (6)$$

$$\sum_{d \in D_k^*} x_{id} \leq \left(\frac{\Phi}{24}\right) y_i \quad \forall i \in I_k \quad (7)$$

$$x_{id} + x_{id'} \leq 1 \quad \forall d \in D_k^*, d' \in Q_k^*(d), i \in I_k \quad (8)$$

$$\sum_{d \in D'} x_{id} = 1 \quad \forall i \in I_k \quad (9)$$

$$y_i \leq y_{i+1} \quad i = 1, \dots, |I_k| - 1 \quad (10)$$

$$x_{id} \in \{0,1\} \quad \forall i \in I_k, d \in D_k^* \quad (11)$$

$$y_i \in \{0,1\} \quad \forall i \in I_k \quad (12)$$

In this mathematical formulation, the objective function (4) minimizes the number of drivers needed at RP k to handle the optimal duties that start and terminate at RP k . Constraint (5) requires assigning one driver to every duty that starts at RP k . Constraint (6) enforces that no driver can handle more than one duty at a time. Constraint (7) requires that a driver can only be assigned to duties if the driver is active. This constraint relates x_{id} and y_i variables, and limits the number of duties assigned to a driver during the planning horizon to satisfy industry regulations (here we assume that the time periods for our model are one hour long). The limitation that no driver can handle two *incompatible duties* – duties that start at the same domicile before the minimum rest period for a driver is completed – is enforced by Constraint (8). Constraint (9) requires that one extended rest period at the domicile has to be assigned to each driver before the end of the planning horizon. This constraint captures the 34-hr restart rule that exists in current hours-of-service regulations for the industry, and can be easily adapted to incorporate the changes that will come into effect in July 2013 according to [15]. Constraint (10) is a symmetry breaking constraint that allows activating a driver only if the immediately higher-numbered driver is active. This constraint helps us to avoid the combinatorial effect of alternative solutions that represent the same driver assignment. Finally, Constraints (11) and (12) require all variables in the model to be binary.

4 Computational Experiments

The building blocks of our proposed technical approach for driver scheduling in TL relay networks were implemented in Python 2.6, and all instances of DSM and DRSCM were solved using CPLEX 12.1 on a Xeon[®] 3.2 GHz workstation with 6 GB of RAM.

4.1 Generation of Instances and Selection of Parameter Values

We generated random problem instances to test the computational performance of our proposed approach presented in Section 3.1 and to analyze the characteristics of the solutions obtained. We generated five instances of 50 node networks by randomly locating uniformly distributed nodes in a region of 600 miles \times 600 miles. This area represents the geographical region covered by a regional network for a major TL carrier. Distances on the arcs were computed using the Euclidean norm as a means to estimate actual over the road distances. For each of our instances, we randomly selected 10% (245) of the origin-destination (O-D) node pairs in the network to have truckload flows. However, since one of our goals is to assess the effect of lane volume on the performance of the approach and the characteristics of the solutions, we varied the actual demand (i.e., number of truckloads shipped) between those selected O-D node pairs. We randomly

generated an integer number uniformly distributed between 10 and 20 for each selected O-D node pair for our *low lane volume* experiments, and between 10 and 40 for our *high lane volume* experiments.

Prior to solving the DSM and DRSCM problems, we solved the relay network design problem to obtain a relay network configuration and a routing for each of the truckloads using the heuristic approach presented in [7]. As a result, the number of hubs varies slightly from one instance to another depending on the optimal number of RPs that was opened as a solution to the relay network design problem. Similarly the number of inter-RP movements (e.g., loads) depends on the routes that are selected for the truckloads in the network. The instances used in our computational experiments considered truckload routes with limitations of 25% circuitry above the shortest path distance between origin and destination of a truckload, and 225 miles and 450 miles for the distances covered by local and lane drivers respectively. Changes to these parameters would likely result in different relay network design configurations and truckload routings that in turn would affect the size of our problem. It is also important to note that the number of duties that are generated not only depends on the limitations that we described in Section 3.2, but it also depends on the design of the relay network (i.e., number and proximity of the RPs). The design of the relay network is affected by the limitations imposed on truckload route circuitry and the distances covered by the drivers. Changes to these parameters will also have an effect on the number of duties that are generated by our proposed approach. This is not explored in the present work.

For the generation of feasible duties to cover the loads in each of our experiments, we considered the following values from hours-of-service regulations [15]: maximum number of driving hours allowed in a day (ω) = 11 hours; minimum number of hours of rest required (τ) = 10 hours; maximum number of hours of rest between two consecutive workdays (τ') = 14 hours; and number of hours of extended rest required (\square) = 34 hours. Although the current regulations do not impose a limitation on the number or frequency of 34-hour restarts, we implemented a portion of the rule that will come into effect in July 2013 by limiting the number of restarts to one in a seven day period (e.g., Constraint (9) of DRSCM). In addition to hours-of-service regulations, we considered a carrier-established limitation on the time away from domicile for drivers of 2 and 3 days to observe the effect of this design parameter on the performance of our modeling approach and the quality of the solutions obtained.

Also, the cost of feasible duties was computed in our duty generation algorithm presented in Section 3.3 by considering a rate of \$1.3 per mile for loaded and empty movements, and a compensation of \$75 per rest period spent away from domicile.

Finally, we analyzed scenarios with planning horizons of 3 and 7 days considering the same freight demand spread over the length of the given planning time period. The purpose was to observe the performance of our approach solving driver scheduling problems across different demand density periods; something that a TL carrier may experience throughout the year.

4.2 Results

Tables 1 and 2 show the results for our *low lane volume* (LLV) experiments when time away from domicile is limited to 3 days. The results in Table 1 correspond to duty selection while driving scheduling results are presented in Table 2.

Table 1: Duty Selection Results for Low Lane Volume and TAFD = 3 days.

<i>PH</i>	<i>Rep.</i>	<i># of Loads</i>	<i># of RPs</i>	<i># of Duties</i>	<i># of Selected Duties</i>	<i>Cost (\$)</i>	<i>Setup Time (secs)</i>	<i>Solution Time (secs)</i>
3	1	3,177	16	199,596	1,445	916,891.72	820.26	111.86
	2	3,786	19	228,517	1,723	1,063,952.61	1,132.95	140.76
	3	3,380	19	144,680	1,593	991,707.71	532.72	74.78
	4	3,589	18	161,776	1,611	1,274,989.30	683.12	70.28
	5	3,565	18	222,915	1,687	948,293.72	1,045.13	67.64
<i>Average</i>		<i>3,499.4</i>	<i>18</i>	<i>191,496.8</i>	<i>1,611.8</i>	<i>1,039,167.01</i>	<i>842.84</i>	<i>93.06</i>
7	1	3,177	16	80,656	1,459	942,824.76	372.85	98.90
	2	3,786	19	99,825	1,729	1,076,513.0	538.40	108.90
	3	3,380	19	63,606	1,602	1,019,958.39	312.39	33.90
	4	3,589	18	75,703	1,644	1,296,142.58	418.53	7.70
	5	3,565	18	91,684	1,699	974,372.81	493.82	66.97
<i>Average</i>		<i>3,499.4</i>	<i>18</i>	<i>82,294.8</i>	<i>1,626.6</i>	<i>1,061,962.31</i>	<i>427.20</i>	<i>63.27</i>

From the results in Table 1, it is clear that spreading the demand over a longer planning horizon (PH) results in a significant reduction in the number of feasible duties in DSM. This is because when loads are spread over a longer planning horizon, it is more challenging to find loads with compatible time windows that can move together in a driver duty. As a result of this, average setup times – the time required to generate the duties using our enumeration algorithm and construct the mathematical model for duty selection – and average solution times are reduced 49.3% and 32% respectively when the average since the number of duties decreases due to the longer planning horizon. As observed in this table, instances with a planning horizon of 3 days were built and solved in less than 22 minutes, while instances with a planning horizon of 7 days were built and solved in less than 11 minutes. However, we can observe that although setup time has a direct relationship to the number of duties in DSM, the solution time varies significantly from one instance to another.

Although problem sizes vary significantly with planning horizon and affect the performance of DSM, the solutions obtained present very similar number of optimal duties selected to cover the loads in the network and no significant difference in the operational costs. However, in order to determine the effect of demand density (i.e., same demand spread over a longer planning horizon) on the cost of scheduling drivers we

need to consider the number of active drivers required to handle the selected duties as determined by the driver scheduling model. In Table 2, it can be observed that the average number of drivers needed for high demand density problems (i.e., 3-day planning horizon problems) is significantly higher both at the domicile level and across the relay network. Thus, driver scheduling and routing can be expected to be more expensive for higher demand density.

Table 2: Driver Scheduling Results for Low Lane Volume and TAFD = 3 days.

<i>PH</i>	<i>Rep.</i>	<i>Avg. # of Vars.</i>	<i>Avg. # of Const.</i>	<i>Avg. # of Drivers per Domicile</i>	<i>Total # of Active Drivers</i>	<i># of Optimal Solutions</i>	<i>Solution Time (secs)</i>	<i>Avg. Opt. Gap</i>	<i>Max. Opt. Gap</i>
3	1	7,257	88,196	48.88	782	15 of 16	28.45	6.25%	6.25%
	2	6,511	68,820	50.05	951	17 of 19	22.56	12.13%	15.17%
	3	6,037	70,430	46.00	874	18 of 19	37.92	40.62%	40.62%
	4	7,728	65,880	61.17	1,101	16 of 18	28.68	19.13%	19.57%
	5	7,401	116,538	48.56	874	16 of 18	51.44	6.92%	10.98%
<i>Average</i>		<i>6,986.8</i>	<i>81,972.8</i>	<i>50.93</i>	<i>896.4</i>		<i>33.81</i>	<i>17.01%</i>	
7	1	5,176	41,713	30.06	481	11 of 16	158.26	49.61%	52.55%
	2	5,140	47,001	29.37	558	14 of 19	175.97	37.27%	71.77%
	3	3,694	27,258	24.26	461	14 of 19	196.60	26.20%	54.23%
	4	5,196	40,499	35.28	635	10 of 18	279.44	22.54%	50.22%
	5	7,128	64,588	29.56	532	12 of 18	54.93	49.39%	56.00%
<i>Average</i>		<i>5,266.8</i>	<i>44,211.8</i>	<i>29.71</i>	<i>533.4</i>		<i>173.04</i>	<i>37.00%</i>	

The results presented in Table 2 for the driver scheduling problem were obtained using the set of optimal duties from DSM as described in Section 3.1. Recall that our approach requires us to provide a set of drivers I_k at each domicile k . We considered total number of truckloads handled at each relay point to obtain an initial estimate for the number of drivers, and used a trial-and-error method to modify this estimate in cases of infeasibility. We also established a time limit for the solution of DRSCM of 15 minutes for problems with a planning horizon of 3 days, and 30 minutes for problems with a planning horizon of 7 days. We report both the number of instances that solved to optimality (# of Optimal Solutions) and the optimality gaps for those instances that stopped after completing the time limit without an optimal solution.

We observed variability from domicile to domicile in the performance of DRSCM at each replication. One of the reasons corresponds to the number of duties that start at a given domicile, with some domiciles having significantly more duties than others. The values presented for number of variables, constraints, and drivers at each domicile in Table 2 are averages across individual domicile problems solved at each replication. As shown in this table, the assignment of duties to drivers cannot be solved to optimality at one or more of the domiciles in each replication when considering a planning horizon of

3 days; however in the column labeled # of Optimal Solutions, we see that the majority of domiciles solved to optimality with only one or two unable to solve to optimality. The column labeled Solution Times presents the average solution times for those problems that were solved to optimality, and the final two columns report the average and worst case optimality gaps for the problems that could not be solved to optimality. As the planning horizon increases to 7 days, we see the tractability issues worsen as each instance has between 5 and 8 domiciles unable to solve to optimality and large optimality gaps at termination for those problems. These problems are more difficult to solve since we are enforcing the 34-hour rule restart established by hours-of-service regulations (i.e., Constraint (9) in DRSCM) which was relaxed for the problems with 3-day planning horizons. In these instances, the number of optimal schedules found is always less than in the case for the shorter planning horizon problems, and the average and maximum optimality gaps are always higher. Solution times for the problems with optimal schedules also increase significantly between the two planning horizons.

Although the difference in the average number of variables is not significant for different levels of freight density (i.e., different planning horizons), the average number of constraints varies significantly with a reduction of 33.7% when the planning horizon increases from 3 to 7 days despite the inclusion of additional constraints to enforce the 34-hour rule as described above. The reason for this decrease is because having the freight spread over a longer planning horizon results in a significant reduction in the number of incompatible duties for every other duty in the model and thus fewer Constraints (8) in the model.

With respect to the solutions found, the average number of drivers per domicile and the total number of active drivers across the relay network are reduced 41.66% and 40.50% respectively when considering the longer planning horizon. Enforcing the 34-hour restart rule is one of the reasons why the reduction in the number of required drivers is not as high as one would expect when more than doubling the length of the planning horizon to serve the same truckload demand. Another reason is because the solutions found for some of the domiciles in each replication that are not optimal. Allowing a longer running time for these instances would likely result in solutions with a lower number of drivers required.

In the following subsections, we consider the effect of further restricting the length of the driver duties and having a higher lane volume in the relay network.

4.2.1 Effect of Time Away From Domicile

We now explore the effect of changing the limitation on TAFD. Increasing or decreasing TAFD will impact which duties are feasible and, consequently, the total number of duties in DSM. In our computational experiments we wanted to quantify this impact as well as to assess the effect on operational cost. This last aspect is important to carriers and researchers who are interested in determining some of the cost efficiency tradeoffs between operating a relay network and using traditional PtP dispatching.

Table 3 shows the results for duty selection considering a limitation of 2 days away from domicile for the drivers. The values between parentheses underneath the average results shown in this table represent the differences with respect to the average values obtained when TAFD = 3 days (Table 1).

Table 3: Duty Selection Results for Low Lane Volume and TAFD = 2 days.

<i>PH</i>	<i>Rep.</i>	<i># of Loads</i>	<i># of RPs</i>	<i># of Duties</i>	<i># of Selected Duties</i>	<i>Cost (\$)</i>	<i>Setup Time (secs)</i>	<i>Solution Time (secs)</i>
3	1	3,177	16	147,605	1,452	927,820.07	565.06	6.95
	2	3,786	19	171,482	1,789	1,111,589.45	808.23	27.67
	3	3,380	19	108,365	1,623	1,021,832.41	400.12	55.01
	4	3,589	18	102,516	1,783	1,393,609.18	458.90	10.46
	5	3,565	18	165,576	1,692	952,632.17	714.73	24.44
	<i>Average</i>	<i>3,499.4</i>	<i>18</i>	<i>139,108.8</i> <i>(-27.36%)</i>	<i>1,667.8</i> <i>(+3.47%)</i>	<i>1,081,496.66</i> <i>(+4.07%)</i>	<i>589.41</i> <i>(-30.07%)</i>	<i>24.91</i> <i>(-73.24%)</i>
7	1	3,177	16	59,258	1,479	956,772.21	338.06	6.68
	2	3,786	19	74,002	1,782	1,112,714.60	491.55	7.89
	3	3,380	19	47,719	1,653	1,047,034.50	291.61	4.00
	4	3,589	18	46,817	1,766	1,368,884.75	373.44	0.94
	5	3,565	18	68,050	1,712	984,297.88	448.30	21.67
	<i>Average</i>	<i>3,499.4</i>	<i>18</i>	<i>59,169.2</i> <i>(-28.10%)</i>	<i>1,678.4</i> <i>(+3.18%)</i>	<i>1,093,940.79</i> <i>(+3.01%)</i>	<i>388.59</i> <i>(-9.04%)</i>	<i>8.23</i> <i>(-86.98%)</i>

As observed in Table 3, a limitation of 2 days for the TAFD results in smaller instance sizes of DSM due to a reduction in the number of duties that are generated. The reduction in problem size is similar for both planning horizons relative to the results obtained when TAFD = 3 days. We also observed that since having duties with up to 2 days away from domicile reduces the number of duties that can cover more than 2 loads; the solutions to DSM include more duties than before. This essentially implies that we need a larger number of shorter duties to cover the same demand. As a result, the cost of routing and rest for the drivers increases as well. However, the increase in number of duties and cost is not very significant and it never exceeds 5% for both planning horizons.

Due to the reduced number of duties being generated in each replication, there is a reduction in setup times. This reduction is more evident for those problems with more duties as it is the case for replications with 3-day planning horizons. For these instances, the average setup time is 30% less than before when longer duties were also generated. However, the biggest effect of reducing TAFD to 2 days is observed in the time that it takes to solve DSM. Problems with a planning horizon of 3 days have a reduction in average solution time that exceeds 73%, while problems with 7-day planning horizons are solved on average more than 86% faster than before. Although the reduction in solution times is very significant, total time required to obtain a solution is still driven by setup time. For 3-day planning horizon problems, instances were created and solved in

less than 14 minutes in the worst case, while problems with planning horizons of 7 days were completed in less than 9 minutes in the worst case.

Table 4 shows the results for driver scheduling when considering TAFD = 2 days and planning horizons of 3 and 7 days.

Table 4: Driver Scheduling Results for Low Lane Volume and TAFD = 2 days.

<i>PH</i>	<i>Rep</i>	<i>Avg. # of Vars.</i>	<i>Avg. # of Const.</i>	<i>Avg. # of Drivers per Domicile</i>	<i>Total # of Active Drivers</i>	<i># of Optimal Solutions</i>	<i>Solution Time (secs)</i>	<i>Avg. Opt. Gap</i>	<i>Max. Opt. Gap</i>
	1	7,229	86,275	49.69	795	15 of 16	30.33	1.16%	1.16%
	2	6,835	75,860	51.37	976	19 of 19	46.06	0.00%	0.00%
3	3	6,225	76,598	45.95	873	19 of 19	68.87	0.00%	0.00%
	4	8,849	96,933	66.00	1,188	18 of 18	68.17	0.00%	0.00%
	5	7,379	109,616	48.72	877	18 of 18	72.42	0.00%	0.00%
	<i>Average</i>	<i>7,303.4 (+4.5%)</i>	<i>89,056.6 (+8.7%)</i>	<i>52.91 (+3.9%)</i>	<i>941.8 (+5.1%)</i>		<i>57.17 (+69.1%)</i>		
	1	7,351	59,274	31.38	502	10 of 16	176.12	53.4%	55.6%
	2	7,638	74,594	32.68	621	14 of 19	152.72	56.1%	78.0%
7	3	3,780	28,742	24.84	472	16 of 19	301.80	37.6%	53.1%
	4	5,358	45,055	32.78	590	10 of 18	350.92	25.1%	62.8%
	5	7,145	63,673	30.50	549	12 of 18	74.33	50.9%	57.1%
	<i>Average</i>	<i>6,254.4 (+18.8%)</i>	<i>54,267.6 (+22.7%)</i>	<i>30.44 (+2.4%)</i>	<i>546.8 (+2.5%)</i>		<i>211.18 (+22.0%)</i>		

As observed in Table 4, since duty selection solutions have more duties when TAFD is limited to 2 days, the average problem size for DRSCM increases as well. As a result of the increase in instance size for DRSCM at each domicile, it takes more time to find optimal solutions. We observe that these problems are significantly easier to solve over shorter planning horizons (3 days) as observed by the number domiciles for each instance that solve to optimality, the time required to obtain the optimal solutions when they can be obtained, and the optimality gap at termination; the performance is significantly worse over longer (7 day) planning horizons.

In terms of the solutions obtained, the average number of drivers per domicile and the total number of active drivers increase as compared to the case with TAFD = 3 days. The average increase for the longer planning horizon problems is less than 2.5%, while it is close to 5% for the 3-day planning horizon problems. A one to one domicile comparison of the number of drivers required when going from 3 to 2 days in TAFD indicates that in most cases the difference is close to the average and there are only a few cases, especially problems without optimal solutions, where the difference is significant. However, this difference never exceeds 23 drivers for problems that are solved to optimality in both cases.

Although the actual difference between the total costs of activating and scheduling drivers when going from 3 to 2 days in TAFD cannot be estimated from our results, we can imply that the effect of different duty lengths seems to mostly depend on the number of drivers required, especially since operational costs are only marginally affected by changes to this parameter. Additional testing with longer duties is necessary to better assess the effect of duty length on the costs associated with operating a relay network and analyzing the tradeoffs with respect to traditional PtP dispatching.

4.2.2 Effect of Lane Volume

To assess the effect of lane volume on the performance of our proposed approach and the characteristics of the solutions found, we considered problem instances with a limitation on TAFD of 2 days and planning horizons of 3 and 7 days. Recall that these *high lane volume* (HLV) experiments consider instances in which we increased the number of truckloads that are shipped between the selected O-D pairs of nodes in the relay networks that we generated. Table 5 shows the DSM results for our HLV experiments. The values between parentheses represent the differences with respect to the average values obtained for the LLV experiments of the same type (Table 3).

Table 5: Duty Selection Results for High Lane Volume and TAFD = 2 days.

<i>PH</i>	<i>Rep.</i>	<i># of Loads</i>	<i># of RPs</i>	<i># of Duties</i>	<i># of Selected Duties</i>	<i>Cost (\$)</i>	<i>Setup Time (secs)</i>	<i>Solution Time (secs)</i>
3	1	5,356	16	491,122	2,438	1,578,260.54	5,167.68	77.58
	2	5,866	19	537,925	2,740	1,695,945.78	6,213.84	206.55
	3	5,545	19	350,981	2,707	1,633,536.45	3,068.92	228.14
	4	5,831	18	299,662	2,889	2,199,388.81	2,864.33	21.87
	5	5,689	18	469,833	2,694	1,727,891.44	4,974.65	124.44
	<i>Average</i>	<i>5,657.4</i> <i>(+61.67%)</i>	<i>18</i>	<i>429,904.6</i> <i>(+209.1%)</i>	<i>2,693.6</i> <i>(+61.51%)</i>	<i>1,727,891.44</i> <i>(+59.77%)</i>	<i>4,457.88</i> <i>(+656.3%)</i>	<i>131.72</i> <i>(+428.8%)</i>
7	1	5,356	16	183,203	2,476	1,590,079.88	1,808.59	21.52
	2	5,866	19	201,226	2,797	1,729,014.99	2,133.18	15.05
	3	5,545	19	142,003	2,697	1,654,905.02	1,437.05	6.55
	4	5,831	18	127,994	2,823	2,140,134.57	1,727.93	12.22
	5	5,689	18	184,806	2,718	1,543,457.79	1,981.64	36.29
	<i>Average</i>	<i>5,657.4</i> <i>(+61.67%)</i>	<i>18</i>	<i>167,846.4</i> <i>(+183.7%)</i>	<i>2,702.2</i> <i>(+61.00%)</i>	<i>1,731,518.45</i> <i>(+58.28%)</i>	<i>1,817.68</i> <i>(+367.8%)</i>	<i>18.33</i> <i>(+122.7%)</i>

As observed in Table 5, the higher volume on the lanes of our instances represents an average increase of approximately 62% more loads in the network. As a result, the number of generated duties increases significantly, especially for problems with 3-day planning horizons where the average problem size of DSM increases more than 3 times as compared to the problem sizes of the LLV experiments. As a consequence, setup and

solution of these higher lane volume problems take considerably more time. This increase in time is very significant for 3-day planning horizon problems which are built and solved approximately 107 minutes in the worst case; this is almost 8 times longer than the LLV experiments of the same type. Similarly, while 7-day planning horizon problems are built and solved in less than 36 minutes, this is still 4 times longer than the time it took to solve similar LLV problems.

Although it is evident that the performance of DSM is challenged by the size of the instances with higher lane volume, the most interesting observation in these experiments relates to the solutions that are obtained. As observed in Table 5, for both planning horizons, the increase in the number of optimal duties selected (approximately 61% more) and the increase in operational costs (approximately 60% more) are very close to the increase in the number of loads in the network (approximately 62%) for the HLV experiments. This is an indication that even though there are non-linear cost structures for the driver duties, there is a direct linear relationship between demand volume and optimal operational costs of routing and rest compensation for the drivers. However, this cost still does not include the activation of drivers at the domiciles in the relay network.

Based on the optimal duties obtained by DSM, the results of DRSCM for the HLV experiments are presented in Table 6.

Table 6: Driver Scheduling Results for High Lane Volume and TAFD = 2 days.

<i>PH</i>	<i>Rep</i>	<i>Avg. # of Vars.</i>	<i>Avg. # of Const.</i>	<i>Avg. # of Drivers per Domicile</i>	<i>Total # of Active Drivers</i>	<i># of Optimal Solutions</i>	<i>Solution Time (secs)</i>	<i>Avg. Opt. Gap</i>	<i>Max. Opt. Gap</i>
	1	19,843	367,575	81.69	1,307	10 of 16	123.73	16.7%	35.1%
	2	17,419	328,073	80.32	1,526	16 of 19	174.56	9.5%	20.6%
3	3	17,215	364,657	74.79	1,421	14 of 19	51.77	21.1%	43.6%
	4	23,691	407,987	105.06	1,891	16 of 18	169.47	5.7%	7.5%
	5	18,530	420,020	78.33	1,410	12 of 18	40.96	20.5%	40.3%
	<i>Average</i>	<i>19,339.6</i> <i>(+164.8%)</i>	<i>377,662.4</i> <i>(+324.1%)</i>	<i>84.04</i> <i>(+58.8%)</i>	<i>1,511.0</i> <i>(+60.4%)</i>		<i>112.10</i> <i>(+96.1%)</i>		
	1	13,667	165,631	58.63	938	7 of 16	323.07	62.3%	70.5%
	2	13,621	185,869	56.95	1,082	7 of 19	164.39	50.6%	76.4%
7	3	11,053	132,073	49.05	932	7 of 19	134.19	53.4%	75.1%
	4	15,629	180,518	71.33	1,284	5 of 18	17.12	56.0%	75.4%
	5	14,225	197,244	57.44	1,034	8 of 18	222.00	58.0%	73.9%
	<i>Average</i>	<i>13,639.0</i> <i>(+118.1%)</i>	<i>172,267.0</i> <i>(+217.4%)</i>	<i>58.68</i> <i>(+92.8%)</i>	<i>1,054.0</i> <i>(+92.8%)</i>		<i>172.15</i> <i>(-18.5%)</i>		

In Table 6, we can observe that since high lane volume increases the number of optimal duties found by DSM as discussed above, the size of the DRSCM problems increases significantly across all replications. For both planning horizons, the average number of variables for each driver scheduling problem at the domiciles more than

doubles the number observed for the LLV case. Consequently, with the increase in problem size, it is more difficult to solve some of the domiciles at each replication. The number of problems that stop before reaching optimality at the pre-specified time limit increases in both planning horizons. Also, for the higher density experiments (i.e., 3-day planning horizon problems), the problems that solve to optimality within the time limit take on average almost twice as much time as the experiments with low lane volume; this is not observed for the problems with planning horizons of 7 days.

Finally, the actual impact of high lane volume on the average number of required drivers is still not clearly defined since for a non-trivial number of domiciles at each replication it is very difficult to obtain the optimal assignment of drivers to duties. However, there is a clear indication that lane volume has a very significant effect on the number of drivers needed, and consequently it will also affect considerably the total costs associated with activating and scheduling drivers in a relay transportation network.

5 Conclusions

In this paper, we presented a new modeling approach for scheduling drivers in a relay network for TL transportation that can be also applied to other hub-and-spoke based trucking networks. The proposed approach decomposes this problem in a series of smaller sub-problems that are solved sequentially. First, an algorithm is used to generate driver duties that start and terminate at a domicile in the relay network and cover one or more loads with pre-defined service requirements. Driver duties are created using templates and their feasibility is checked to ensure adherence to hour-of-service regulations, a carrier-based limitation on the time allowed away from domicile for the drivers and the service requirements for the loads. These feasible driver duties are composite variables that are then used in a mathematical formulation (DSM) that selects those that will satisfy demand while minimizing the costs of routing trucks and compensating drivers that spend rest periods away from their domiciles. The assignment of drivers to selected duties (and consequently the creation of their driving schedules for a given planning horizon) is made by solving another mathematical problem (DRSCM) for each individual domicile in the relay network. Decomposing the driver scheduling problem at the domicile level does not come at the expense of a loss of optimality since the solution to DSM provides a unique assignment of loads to driver duties.

After completing computational experiments of our proposed approach, we have determined that the number of variables (i.e., duties) in the model significantly affects the performance of DSM. Our experimentation shows that problem size is affected by all of these three factors: different freight densities (i.e., same demand over different planning horizons), lane volumes (i.e., more truckloads shipped in the same lanes over a given planning horizon), and limitations on the time away from domicile for the drivers. In general, bigger problem instances of DSM take significantly more time to generate feasible duties and to solve. Specifically, freight density and lane volume have a substantial effect on the performance of DSM, whereas changes in TAFD have a still

sizeable but less significant influence in the time required to generate and select feasible duties that are cost effective. Interestingly, out of the three factors mentioned above, only lane volume seems to have a significant effect on the solutions obtained in terms of number of optimal duties needed to satisfy the demand and the operational costs. In this case, there seems to be a direct linear relationship between changes in the number of loads in the network and the subsequent changes in the optimal solutions obtained with DSM. The differences between the solutions found for different freight densities and duty lengths allowed are only marginal.

Similarly, the performance of DRSCM is affected by instance size. The greater the number of duties obtained to the DSM problem, the larger the size of problem instances of DRSCM we observe. However, some instances of these problems are more difficult to solve than others of similar size. As a result, there is significant variability in the performance of the driver scheduling problem across different domiciles in a single replication. Again, high lane volume problems are much more difficult to solve than their low lane volume counterparts. In each of our replications, we observed that some of the schedules for certain domiciles are not optimal given that the solution of DRSCM was stopped after a pre-specified time limit has passed. In fact, the average number of optimal schedules that are obtained in each replication is reduced for longer planning horizons, longer duties for the drivers and especially higher lane volumes.

Finally, there is evidence that the total costs of activating and scheduling drivers at the RPs across the relay network depend mostly on lane volume and freight density. Although we only observed a very modest increase in total costs when changing the limitation on TAFD from 3 to 2 days, it is still necessary to consider longer duties to fully understand the effect of this parameter on the expenses that carriers can expect when operating a relay network. This analysis will also help us to assess the cost efficiency tradeoffs associated with network-based configurations as compared to current PtP dispatching for the truckloads. Other areas for future research are identified in the next section.

6 Future Research

We were able to identify several areas for future work associated with driver scheduling in TL relay networks and other transportation systems based on hub-and-spoke networks.

First, a column generation approach using a multilabel shortest path algorithm to generate duties can be implemented to solve DSM given the large number of variables needed. This approach will allow us to consider duties with longer horizons for TAFD and help us to avoid tractability issues due to problem size.

Similarly, there is a need to explore alternative formulations for the driver scheduling problem at each domicile (DRSCM) to improve on the performance of the model presented in this research. An interesting aspect related to this problem relates to the estimation of the number of drivers that should be used at each RP to obtain a feasible

solution. Some techniques can be explored to determine initial estimates that will facilitate obtaining optimal solutions for DRSCM in an efficient way.

Another area of future work relates to the development of an integrated modeling approach for relay network design and driver scheduling that will make these decisions simultaneously. An integrated approach has the potential to provide better solutions than the decomposition approach presented in this paper since it would directly capture the interactions that exist between these two problems.

Also, the present work can be extended to consider different types of drivers that exist in TL and LTL operations such as those who bid for certain loads. In the present work, any driver can move any load as long as he or she returns home within a pre-specified time limit. Similarly, our approach can be extended to consider union rules in addition to the hours-of-service regulations when determining the feasibility of driver duties.

It would also be interesting to understand whether balancing equipment at the relay points as required by the relay network design problem is an adequate surrogate to enforcing driver balance. We are interested in determining what are the incremental costs related to driver scheduling in addition to those incurred to balance the equipment throughout the network.

Finally, to further support the case for the use of relay networks in practice, it is necessary to understand and quantify the positive service impact of operating a relay network for freight movement as opposed to traditional point-to-point dispatching. This could be a significant step towards the acceptance of this type of systems and highlight the importance of collaborative logistics efforts like the Physical Internet initiative.

Acknowledgements

We would like to acknowledge the work of Sergio Maldonado who helped us with the generation of problem instances for our computational experiments.

References

- [1] Taha, T. T. and Taylor, G. D., “An Integrated Modeling Framework for Evaluating Hub-and-Spoke Networks in Truckload Trucking,” *Logistics and Transportation Review*, 30, 2, 141-166 (1994).
- [2] American Trucking Associations, “Trucking Activity Report,” *Technical Report* (2010).
- [3] Taylor, G. D., Meinert, T. S., Killian, R. C. and Whicker, G. L., “Development and Analysis of Alternative Dispatching Methods in Truckload Trucking,” *Transportation Research: Part E*, 35, 191-205 (1999).

- [4] Taylor, G. D., Whicker, G. L. and Usher, J. S., "Multi-zone Dispatching in Truckload Trucking," *Transportation Research: Part E*, 37, 375-390 (2001).
- [5] Üster, H. and Maheshwari, N., "Strategic Network Design for Multi-zone Truckload Shipments," *IIE Transactions*, 39, 177-189 (2007).
- [6] Üster, H. and Kewcharoenwong, P., "Strategic Design and Analysis of a Relay Network in Truckload Transportation," *Transportation Science*, 45, 4, 505-523 (2011).
- [7] Vergara, H. A. and Root, S., "Composite Variable Model for Truckload Relay Network Design," *University of Arkansas, Department of Industrial Engineering – White Paper*, (2012).
- [8] American Transportation Research Institute, "Critical Issues in the Trucking Industry – 2011," *Technical Report*, (2011).
- [9] Morris, F., "Thousands of Trucking Jobs, But Few Take the Wheel," *National Public Radio*, <http://www.npr.org/2011/10/31/141791952/thousands-of-trucking-jobs-but-few-take-the-wheel> (October 31 2011).
- [10] Montreuil, B., "Toward a Physical Internet: Meeting the Global Logistics Sustainability Grand Challenge," *Logistics Research*, 3, 71-87 (2011).
- [11] Erera, A., Karacik, B. and Savelsbergh, M., "A Dynamic Driver Management Scheme for Less-Than-Truckload Carriers," *Computers & Operations Research*, 35, 3397-3411 (2008).
- [12] Erera, A., Hewitt, M., Karacik, B. and Savelsbergh, M., "Locating drivers in a Trucking Terminal Network," *Transportation Research: Part E*, 45, 988-1005 (2009).
- [13] Erera, A., Hewitt, M., Savelsbergh, M. and Zhang, Y., "Creating Schedules and Computing Operating Costs for LTL Load Plans," to appear in *Computers & Operations Research*, doi:10.1016/j.cor.2011.10.001 (2011).
- [14] Goel, A. and Kok, L., "Truck Driver Scheduling in the United States," to appear in *Transportation Science*, doi:10.1287/trsc.1110.0382 (2011).
- [15] U.S. Department of Transportation – Federal Motor Carrier Safety Administration (FMCSA), "Hours of Service of Drivers," <http://www.fmcsa.dot.gov/rules-regulations/topics/hos/index.htm> (2011).