Behavioral Logistics and Fatigue Management in Vehicle Routing and Scheduling Problems

Zachary E. Bowden

Dissertation submitted to the faculty of the Virginia Polytechnic Institute and State University in partial fulfillment of the requirements for the degree of

Doctor of Philosophy

In

Business Information Technology

Cliff T. Ragsdale, Chair

Debbie F. Cook

Jason K. Deane

Lance A. Matheson

Onur Seref

March 16, 2016

Blacksburg, Virginia

Keywords: Vehicle Routing Problem, Behavioral Logistics, Alertness, Fatigue Management

Behavioral Logistics and Fatigue Management in Vehicle Routing and Scheduling Problems

Zachary E. Bowden

Abstract

The vehicle routing problem (VRP), is a classic optimization problem that aims to determine the optimal set of routes for a fleet of vehicles to meet the demands of a set of customers. The VRP has been studied for many decades and as such, there are many variants and extensions to the original problem. The research presented here focuses on two different types of vehicle routing and scheduling planning problems: car shipping and fatigue-aware scheduling. In addition to modeling and solving the car shipping problem, this research presents a novel way for ways in which drivers can describe their route preferences in a decision support system. This work also introduces the first fatigue-aware vehicle scheduling problem called the Truck Driver Scheduling Problem with Fatigue Management (TDSPFM). The TDSPFM is utilized to produce schedules that keep the drivers more alert than existing commercial vehicle regulations. Finally, this work analyzes the effect of the starting alertness level on driver alertness for the remainder of the work week and examines a critical shortcoming in existing regulations.

Acknowledgements

First, I acknowledge and thank God for life, my family, the people supporting me, and the opportunity to work on this dissertation.

I am thankful for my dissertation chairman, Dr. Cliff Ragsdale who has been to me an exceptional role model and mentor, not only in academia but also in life. I sincerely appreciate the guidance, ideas, and support of my committee members: Drs. Debbie Cook, Jason Deane, Lance Matheson, and Onur Seref. This work was made better because of your ideas and input.

Thank you to Clark Gaylord for being a supporting and accommodating boss. I would also like to thank Rich Hanowski for his time and expertise. Likewise, Tom Dingus, Jon Hankey, and the VTTI community: I am thankful for being able to work in an environment that encourages pursuits such as this.

Finally, I thank my wife Lori and our children Katie, Susanna, Jake, and Owen. You guys have been so good to me during this process. I am thankful for each of you and I love you dearly. I thank you for your encouragement, support, and for keeping me focused on what is really important.

Table of Contents

List of Tables

List of Figures

Chapter 1: Introduction and Literature Review

1.0 Introduction

The vehicle routing problem (VRP), is a classic optimization problem that aims to determine the optimal set of routes for a fleet of vehicles to meet the demands of a set of customers. The VRP has been studied in the Operations Research literature since its introduction in 1959 (Dantzig & Ramser, 1959). Since 1959, VRP research has both enhanced our ability to model and solve this type of problem as well as broadened the definition of what constitutes a VRP through extensions to the original model.

Golden, Raghavan, and Wasail pose that "Vehicle routing may be the single biggest success story in operations research" (2008). This is comforting considering that "VRP is one of the most studied combinatorial optimization problems" (Baldacci, Battarra, & Vigo, 2008).

Since the VRP has been studied for over a half of a century, the extensions to the original problem have been numerous and broad. They range from adding constraints to reflect real-world conditions such as adding time windows to customer sites, to incorporating other types of optimization problems such as driver scheduling. Today, much VRP research is focused on more complex variants of the traditional VRP considering constraints and problem variations motivated by real-world conditions. Toth and Vigo (2002) offer an excellent overview of the types of problem variants and solution methodologies covering the first 40 years of VRP research.

In the traditional VRP the objective function is to minimize delivery cost; often this is modeled by minimizing the duration of the routes. However variations exist with many different objective functions such as maximizing profit or minimizing cargo spoilage. In the research presented here we will utilize two different objective functions depending on the problem we are

working on. We will use the traditional cost minimization objective function as well as maximizing profit by choosing the optimal set of jobs to fulfill. Profit maximization is often the focus when the set of jobs contains optional jobs, which will be the case for some of our research presented below.

The research presented here focuses on two different types of vehicle routing and scheduling planning problems: car shipping and fatigue-aware scheduling. In the car shipping problem, the objective will be to maximize profit under complex constraints. In the fatigue-aware scheduling problem, the objective is to minimize duration while maintaining an acceptable alertness level.

The research presented below addresses three primary questions using the two problems described above:

- How can we select the optimal set of shipping requests to be most profitable in a complex shipping system while also handling behavior specific preferences of the driver?
- How can we plan schedules to ensure that drivers maintain a safe level of alertness while still meeting delivery deadlines, abiding by regulations regarding driving times, and are not substantially more costly than existing routes that ignore fatigue and alertness of drivers?
- Do existing hours of service regulations ensure that drivers maintain an acceptable level of alertness throughout their trip or are there potential gaps in the regulations (such as beginning the work week fatigued) that may contribute to lower alertness levels?

These questions will be explored below where we will build upon the volumes of existing vehicle routing and scheduling research that currently exists in the literature. To accomplish this, we will introduce new models including the first fatigue-aware vehicle scheduling problem. Additionally, we will present a novel idea for the way in which drivers can describe their route preferences in a decision support system.

1.1 Vehicle Routing and Scheduling Problem Varieties

Whereas Toth and Vigo (2002) looked to enumerate nearly the entire range of VRP variants, this research will focus on a smaller subset of the literature. We will introduce two new variants of the VRP built upon these existing variations:

- Capacitated Vehicle Routing Problem (CVRP)

- Vehicle Routing with Pickups and Deliveries (VRPPD)

- Pickup and Delivery Selection Problem (PDSP)

- General Pickup and Delivery Problem (GPDP)

- Truck Driver Scheduling Problem (TDSP)

The Capacitated Vehicle Routing Problem (CVRP)

The CVRP is the problem of selecting routes for a fleet of trucks that simultaneously meet the demand of customers while abiding by the capacity of the trucks. Trucks in a CVRP are loaded at a depot location and then travel to the customer sites. Typically, other constraints are also added such as time windows, truck types, and possibly optional deliveries. The typical objective function of the CVRP is to minimize the cost of serving the selected routes. The CVRP is a well-studied problem in the literature that has recently been studied with respect to exact algorithms (Baldacci, Mingozzi, & Roberti, 2012), demand uncertainty (Gounaris, Wiesemann,

4

& Floudas, 2013), and optimizing with respect to specifically minimizing fuel consumption as opposed to cost or distance (Xiao, Zhao, Kaku, & Xu, 2012). Additional overview references pertaining to the CVRP are available in (Golden et al., 2008; Toth & Vigo, 2002).

The Vehicle Routing Problem with Pickups and Deliveries (VRPPD)

In the VRPPD, the goal is to select an optimal set of routes to meet the demand of customers at minimal cost. However in the VRPPD, each transportation request has an origin and a destination rather than the origin always being a depot location. Typically, models assume that all requests must be fulfilled and ignore situations where job selection would be required (as required in this research). Frequently, loads being shipped in the VRPPD are handled as commodities in that there is no specific destination for an item once it is picked up, it just will need to be delivered to some location that has a need for the item (i.e. a delivery node). In the VRPPD literature, we have a specific interest in the Single Vehicle One-to-One Pickup and Delivery Problem (SVPDP) (Cordeau, Laporte, & Ropke, 2008). The VRPPD is a well-studied problem with over 30 years of published research support (Berbeglia, Cordeau, Gribkovskaia, & Laporte, 2007). As such, many surveys and generalized discussions exist including (Savelsbergh & Sol, 1995; Toth & Vigo, 2002).

The Pickup and Delivery Selection Problem (PDSP)

As the name implies, the PDSP is focused on the optimal selection of transportation requests. In contrast to the CVRP or VRPPD, the typical objective function in the PDSP is to maximize profit by selecting the most profitable jobs. The general focus of the PDSP is on selection of jobs as opposed to routing. However, optimal route selection can be included in the model if such an objective is desired (Maes, Caris, Ramaekers, Janssens, & Bellemans, 2012;

Savelsbergh & Sol, 1995; Schönberger, Kopfer, & Mattfeld, 2003). In the reviewed literature discussing problems modeled as a PDSP or a variant of the PDSP, researchers focus their solution efforts on large problems and heuristic solutions.

The General Pickup and Delivery Problem (GPDP)

Savelsbergh and Sol discuss the characteristics of a pickup and delivery problem that differs from a typical VRP (Savelsbergh & Sol, 1995). In addition, they construct a flexible model that can handle the typical instances of pickup and delivery problems and coin the term *General Pickup and Delivery Problem*. In chapter 2 of this dissertation we will build upon the GPDP model and add missing components motivated from the three VRP variants discussed above that are required to accurately model our problem of interest.

The Truck Driver Scheduling Problem (TDSP)

In the TDSP, the focus is on the schedule of deliveries and other stops (such as rest breaks) as opposed to the routes to be taken. In the context of the TDSP, the sequence of visits is known ahead of time and the goal is to set the optimal schedule of deliveries, rest breaks, and other stops such that the total duration is minimized while no scheduling constraints are violated. We are particularly interested in scheduling constraints related to regulations of truck drivers known as the Hours of Service (HOS) restrictions. The TDSP has been studied most extensively by Goel in (Goel, 2010, 2012a, 2012b) as well as Archetti and Savelsbergh (2009). In chapter 3 we will built upon the existing TDSP to introduce a fatigue-aware TDSP that will allow us to measure and control fatigue levels in a schedule.

1.1.1 Objective: Maximize Profit

In a VRP with a profit maximization objective function, the broad goal is to select as many profitable jobs as possible with some finite set of resources. In the case we consider (and as is often the case) the limitations are in capacity and time. We will also be studying direct shipping procurement which implies a pickup and delivery associated with each job. A job can be optional; however, once the pickup has been made the respective delivery must also be made within the time window defined.

In our case, we will be studying the relatively new market of direct consumer procurement of a shipping provider for large items (such as a car). A lucrative market exists for shipping providers to meet the increased demand of shipping these large items. We are therefore motivated to study the way in which shipping providers should select jobs and associated routes to fulfill customers' demands in order to maximize profit.

We were inspired to consider this problem by the relatively new phenomenon of online car shopping on websites such as eBay Motors. Thousands of vehicles are sold on eBay Motors every day and an increasing number of these sales are taking place with significant distances between the buyer and seller. In such cases, the buyer has two options to receive the vehicle: 1) travel to the vehicle and drive it home, or, 2) have the vehicle shipped to them. We focus on the latter option which has seen a growth in popularity in recent years (uShip.com, 2015).

Five or ten years ago, this type of purchase was fairly rare for the general consumer; however, today it occurs hundreds or even thousands of times per day. eBay Motors alone has facilitated the sale of over 5.2 million vehicles online (eBay Inc., 2015). During the 3rd quarter of 2012, which was the last time eBay released the necessary detailed sales information, the Gross

Merchandise Value (GMV) of vehicle sold on eBay Motors was approximately $2 billion. Of the sales in Q3 of 2012 as well as in Q3 of 2014, 77% were interstate transactions increasing the likelihood of the need to ship the newly purchased vehicle (eBay Inc., 2012, 2015).

Broadly speaking, we attempt to determine the optimal set of arcs in a road network for a delivery vehicle to traverse. As such, this type of problem is a variant of the vehicle routing problem (VRP), which as mentioned previously is a classic optimization problem. Since the VRP's introduction in 1959, VRP research has both enhanced our ability to model and solve this type of problem as well as broadened the definition of what constitutes a VRP through extensions to the original model. Online shipping procurement offers a new opportunity to broaden the VRP through the unique aspects of the problem that must be considered and modeled.

The decision problem we address involves the profit maximizing selection of the jobs to accept and delivery route to take given the capacity constraints on the delivery vehicle, the pickup and delivery points of the non-fungible items associated with each job, the pickup and delivery time windows associated with each job, and a user-defined backtracking limitation which will be introduced and discussed in detail in chapter 2. We have been unable to find published research that addresses this specific combination of issues; however we will present a review of a sizable array of research that has investigated various individual components that make up our problem.

1.1.2 Objective: Remain Alert

Driver fatigue has been empirically identified as a major factor in vehicle crashes (Chen & Xie, 2014). Studies by the National Highway Traffic Safety Administration (NHTSA) show

that 100,000 police-reported crashes, 71,000 injuries, and 1,550 deaths occur due to drowsy driving each year in the United States (Rau 2005). Bowman et al. (2012) conclude that driver fatigue is the probable cause of 30% of crashes. According to the National Sleep Foundation, these statistics "significantly underestimate the problem" (National Sleep Foundation 2007) and there exist naturalistic driving studies to support the claim of underestimation (Dingus et al. 2006; Klauer et al. 2006).

For commercial trucks, many countries (and the European Union) have hours of service (HOS) regulations that attempt to control the amount of fatigue a truck driver experiences while driving (Cappuccio, Miller, & Lockley, 2010, pp. 424–429; Rancourt, Cordeau, & Laporte, 2012; "Summary of Hours of Service Regulations," 2013). When these regulations are modified, studies are published attempting to quantify the effect of the changes on safety outcomes (Blanco et al., 2011; Hanowski, Hickman, Olson, & Bocanegra, 2009). These regulations attempt to balance the safety benefits with the added costs that will be incurred by compliance. Quantifying these benefits, costs, and other factors introduced with regulatory compliance is an area of active research which seems to point to significant room for improvement (Heaton, 2005; Jensen & Dahl, 2009; Kemp, Kopp, & Kemp, 2013).

Mathematical models and solutions for the Truck Driver Scheduling Problem (TDSP) and Vehicle Routing Problem (VRP) incorporating HOS constraints have been introduced in the literature (Archetti & Savelsbergh, 2009; Goel, 2012a, 2014; Xu, Chen, Rajagopal, & Arunapuram, 2003). Goel and Vidal (2014) model and compare several HOS constraints from both a cost and risk standpoint. However, we have been unable to find evidence of a model that incorporates fatigue considerations into the model itself.

Mathematical models for predicting alertness have existed at least since Borbely's two-process model was published in 1982 (Borbely, 1982). Since then several additional models have been published, many of which are derived from Borbely's initial work. See Mallis et al. (2004), Gundel et al. (2007), and Dawson et al. (2011) for a review of these models. The accuracy of these models has been shown to be similar (Van Dongen, 2004). Specifically with respect to vehicle crashes, it has been shown that one can reasonably predict crashes with a mathematical model that uses a sleep/wake predictor that is based on the Three Process Model of Alertness (Åkerstedt, Connor, Gray, & Kecklund, 2008). The Three Process Model of Alertness (TPMA) is itself an extension of Borbely's original model and we will leverage the most recent implementation and parameter values for the TPMA presented in Ingre et al. (2014).

We introduce the Truck Driver Scheduling Problem with Fatigue Monitoring (TDSPFM) by incorporating the Three Process Model of Alertness. This allows us to choose a schedule that keeps the driver alert and thereby reduce the likelihood of being involved in a fatigue related crash. We assume that the most recent United States HOS regulations (Goel, 2014) must be obeyed and include constraints to represent these restrictions. This will allow us to calculate the alertness of a schedule that minimizes its duration while abiding by the HOS regulations. We will use this as a benchmark and compare schedules where we maintain certain minimum levels of alertness.

We will then continue this investigation and study the effect of different starting alertness levels. We noticed that in all HOS regulations we reviewed, none have provisions on how rested a driver needs to be when they begin their work. This means that someone could be awake for more than 24 hours straight and begin driving a fully loaded tractor trailer for another 8 hours before being required to take their first break and still be in compliance with US HOS

regulations. A publicized example of this occurred on June 7, 2014 in the fatal crash involving comedians Tracy Morgan and James McNair (Hanna & Marsh, 2015). Obviously if a driver has been awake for an extended period of time then his/her alertness at that particular point is likely to be poor. However, we extend that idea to look at how different levels of fatigue at the beginning of the route affect the alertness of the driver for the remainder of the week.

1.2 Future Research

Since we are introducing the first TDSPFM model in the literature, there will be significant opportunity to validate and expand the model. In chapter 4 we will show how the TDSPFM performs on an existing real-world truck driver schedule extracted from an actual log book. We can further validate our model by using data sets obtained as part of Federal Motor Carrier Safety Administration studies such as the "Commercial Motor Vehicle Driver Restart Study" (FMCSA, 2015). Additionally we could incorporate a study of driver schedule plans as part of a naturalistic driving study at the Virginia Tech Transportation Institute (VTTI).

There are two primary areas we can look at for expansion of the TDSPFM model: sleep variability and driving specific alertness modifications. We can look at sleep variability from both the standpoint of individual variations and specific rest period variations. In our current implementations presented below, we look at the typical/average sleep functions. The TPMA however, can easily support user specific functions. For instance, certain people may recover during sleep at different rates or be more affected by being awake and needing to work at 3am. Additionally, our models assume non-stochastic sleep times during rest periods. However, practically speaking the amount of time that you actually sleep on a given night is dependent

upon many factors (noise levels, comfort, temperature, etc.). We could incorporate a stochastic element into the model to represent how these unknown factors affect sleep duration.

As a planning problem, we should look for areas where a decision support system (DSS) could aid in selecting safer routes for drivers. This could allow for selection of parameter values such as a minimum acceptable alertness level, or it could also alert schedulers to specific points along the route where alertness will be near an unsafe level. This could allow some specific action to be taken to, for example, check in with the driver around that time or take extra precautions to increase the chances that the driver got a good night's rest the night prior. Other extensions exist such as allowing for different alertness levels depending on the cargo being transported or the areas through which the route passes.

References

Åkerstedt, T., Connor, J., Gray, A., & Kecklund, G. (2008). Predicting road crashes from a mathematical model of alertness regulation—The Sleep/Wake Predictor. *Accident Analysis & Prevention*, *40*(4), 1480–1485. http://doi.org/10.1016/j.aap.2008.03.016

Archetti, C., & Savelsbergh, M. (2009). The Trip Scheduling Problem. *Transportation Science*, *43*(4), 417–431. http://doi.org/10.1287/trsc.1090.0278

Baldacci, R., Battarra, M., & Vigo, D. (2008). Routing a Heterogeneous Fleet of Vehicles. In B. Golden, S. Raghavan, & E. Wasil (Eds.), *The Vehicle Routing Problem: Latest Advances and New Challenges* (pp. 3–27). Springer US. Retrieved from http://link.springer.com/chapter/10.1007/978-0-387-77778-8_1

Baldacci, R., Mingozzi, A., & Roberti, R. (2012). Recent exact algorithms for solving the vehicle routing problem under capacity and time window constraints. *European Journal of Operational Research*, *218*(1), 1–6. http://doi.org/10.1016/j.ejor.2011.07.037

Berbeglia, G., Cordeau, J.-F., Gribkovskaia, I., & Laporte, G. (2007). Static pickup and delivery problems: a classification scheme and survey. *TOP*, *15*(1), 1–31. http://doi.org/10.1007/s11750-007-0009-0

Blanco, M., Hanowski, R. J., Olson, R. L., Morgan, J. F., Soccolich, S. A., Wu, S.-C., & Guo, F. (2011). The Impact of Driving, Non-Driving Work, and Rest Breaks on Driving Performance in Commercial Motor Vehicle Operations. Retrieved from http://trid.trb.org/view.aspx?id=1108816

Borbely, A. A. (1982). A Two Process Model of Sleep Regulation. *Human Neurobiol.*, *1*, 195–204.

Bowman, D. S., Schaudt, W. A., & Hanowski, D. R. J. (2012). Advances in Drowsy Driver Assistance Systems Through Data Fusion. In A. Eskandarian (Ed.), *Handbook of Intelligent Vehicles* (pp. 895–912). Springer London. Retrieved from http://link.springer.com/referenceworkentry/10.1007/978-0-85729-085-4_34

Cappuccio, F., Miller, D. M. A., & Lockley, S. W. (2010). *Sleep, Health, and Society: From Aetiology to Public Health*. Oxford University Press.

Chen, C., & Xie, Y. (2014). The impacts of multiple rest-break periods on commercial truck driver's crash risk. *Journal of Safety Research*, *48*, 87–93. http://doi.org/10.1016/j.jsr.2013.12.003

Cordeau, J.-F., Laporte, G., & Ropke, S. (2008). Recent Models and Algorithms for One-to-One Pickup and Delivery Problems. In B. Golden, S. Raghavan, & E. Wasil (Eds.), *The Vehicle Routing Problem: Latest Advances and New Challenges* (pp. 327–357). Springer US. Retrieved from http://link.springer.com/chapter/10.1007/978-0-387-77778-8_15

Dantzig, G. B., & Ramser, J. H. (1959). The Truck Dispatching Problem. *Management Science*, *6*(1), 80–91. http://doi.org/10.1287/mnsc.6.1.80

Dawson, D., Ian Noy, Y., Härmä, M., Åkerstedt, T., & Belenky, G. (2011). Modelling fatigue and the use of fatigue models in work settings. *Accident Analysis & Prevention*, *43*(2), 549–564. http://doi.org/10.1016/j.aap.2009.12.030

eBay Inc. (2012, October 16). *eBay Motors Fast Facts Q3 2012*. Retrieved from http://legacy.ebayinc.com/assets/pdf/fact_sheet/eBay_Motors_Fact_Sheet_Q32012.pdf

eBay Inc. (2015, January 15). *eBay Marketplaces Fast Facts At-A-Glance (Q3 2014)*. Retrieved from http://www.ebayinc.com/sites/default/files/MP_FastFacts_Q32014_0.pdf

FMCSA. (2015, April 28). Commercial Motor Vehicle Driver Restart Study [Text]. Retrieved September 1, 2015, from http://www.fmcsa.dot.gov/safety/research-and-analysis/commercial-motor-vehicle-driver-restart-study

Goel, A. (2010). Truck Driver Scheduling in the European Union. *Transportation Science*, *44*(4), 429–441. http://doi.org/10.1287/trsc.1100.0330

Goel, A. (2012a). The Canadian minimum duration truck driver scheduling problem. *Computers & Operations Research*, *39*(10), 2359–2367. http://doi.org/10.1016/j.cor.2011.12.016

Goel, A. (2012b). *The Minimum Duration Truck Driver Scheduling Problem* (SSRN Scholarly Paper No. ID 1798569). Rochester, NY: Social Science Research Network. Retrieved from http://papers.ssrn.com/abstract=1798569

Goel, A. (2014). Hours of service regulations in the United States and the 2013 rule change. *Transport Policy*, *33*, 48–55. http://doi.org/10.1016/j.tranpol.2014.02.005

Goel, A., & Vidal, T. (2014). Hours of Service Regulations in Road Freight Transport: An Optimization-Based International Assessment. *Transportation Science*, *48*(3), 391–412. http://doi.org/10.1287/trsc.2013.0477

Golden, B. L., Raghavan, S., & Wasil, E. A. (2008). *The Vehicle Routing Problem: Latest Advances and New Challenges*. Springer.

Gounaris, C. E., Wiesemann, W., & Floudas, C. A. (2013). The Robust Capacitated Vehicle Routing Problem Under Demand Uncertainty. *Operations Research*. http://doi.org/10.1287/opre.1120.1136

Gundel, A., Marsalek, K., & Thoren, C. ten. (2007). A critical review of existing mathematical models for alertness. *Somnologie - Schlafforschung Und Schlafmedizin*, *11*(3), 148–156. http://doi.org/10.1007/s11818-007-0312-x

Hanna, J., & Marsh, R. (2015, August 11). NTSB: Trucker's fatigue likely led to Tracy Morgan crash - CNN.com. *CNN*. Retrieved from http://www.cnn.com/2015/08/11/us/ntsb-tracy-morgan-crash/index.html

Hanowski, R. J., Hickman, J. S., Olson, R. L., & Bocanegra, J. (2009). Evaluating the 2003 revised hours-of-service regulations for truck drivers: The impact of time-on-task on critical incident risk. *Accident Analysis & Prevention*, *41*(2), 268–275. http://doi.org/10.1016/j.aap.2008.11.007

Heaton, K. (2005). Truck Driver Hours of Service Regulations: The Collision of Policy and Public Health. *Policy, Politics, & Nursing Practice*, *6*(4), 277–284. http://doi.org/10.1177/1527154405282841

Ingre, M., Van Leeuwen, W., Klemets, T., Ullvetter, C., Hough, S., Kecklund, G., … Åkerstedt, T. (2014). Validating and Extending the Three Process Model of Alertness in Airline Operations. *PLoS ONE*, *9*(10), e108679. http://doi.org/10.1371/journal.pone.0108679

Jensen, A., & Dahl, S. (2009). Truck drivers hours-of-service regulations and occupational health. *Work: A Journal of Prevention, Assessment and Rehabilitation*, *33*(3), 363–368. http://doi.org/10.3233/WOR-2009-0884

Kemp, E., Kopp, S. W., & Kemp, E. (2013). Six days on the road: Will I make it home safely tonight? Examining attitudes toward commercial transportation regulation and safety. *The International Journal of Logistics Management*, *24*(2), 210–229.

Maes, T., Caris, A., Ramaekers, K., Janssens, G. K., & Bellemans, T. (2012). Pickup and delivery selection problem. Retrieved from https://uhdspace.uhasselt.be/dspace/handle/1942/14365

Mallis, M. M., Mejdal, S., Nguyen, T. T., & Dinges, D. F. (2004). Summary of the Key Features of Seven Biomathematical Models of Human Fatigue and Performance. *Aviation, Space, and Environmental Medicine*, *75*(3), A4–A14.

Rancourt, M.-E., Cordeau, J.-F., & Laporte, G. (2012). Long-Haul Vehicle Routing and Scheduling with Working Hour Rules. *Transportation Science*, *47*(1), 81–107. http://doi.org/10.1287/trsc.1120.0417

Savelsbergh, M. W. P., & Sol, M. (1995). The General Pickup and Delivery Problem. *Transportation Science*, *29*(1), 17–29. http://doi.org/10.1287/trsc.29.1.17

Schönberger, J., Kopfer, H., & Mattfeld, D. C. (2003). A Combined Approach to Solve the Pickup and Delivery Selection Problem. In P. D. U. Leopold-Wildburger, P. D. F. Rendl, & P. D. G. Wäscher (Eds.), *Operations Research Proceedings 2002* (pp. 150–155). Springer Berlin Heidelberg. Retrieved from http://link.springer.com/chapter/10.1007/978-3-642-55537-4_24

Summary of Hours of Service Regulations. (2013, December 30). [Text]. Retrieved April 22, 2015, from http://www.fmcsa.dot.gov/regulations/hours-service/summary-hours-service-regulations

Toth, P., & Vigo, D. (2002). *The Vehicle Routing Problem*. Society for Industrial and Applied Mathematics. Retrieved from http://books.google.com/books?id=TeMgA5S74skC

uShip.com. (2015, September 1). *uShip.com - About*. Retrieved from
http://www.uship.com/about/

Van Dongen, H. P. A. (2004). Comparison of Mathematical Model Predictions to Experimental
Data of Fatigue and Performance. *Aviation, Space, and Environmental Medicine*, *75*(3),
A15–A36.

Xiao, Y., Zhao, Q., Kaku, I., & Xu, Y. (2012). Development of a fuel consumption optimization
model for the capacitated vehicle routing problem. *Computers & Operations Research*,
39(7), 1419–1431. http://doi.org/10.1016/j.cor.2011.08.013

Xu, H., Chen, Z.-L., Rajagopal, S., & Arunapuram, S. (2003). Solving a Practical Pickup and
Delivery Problem. *Transportation Science*, *37*(3), 347–364.
http://doi.org/10.1287/trsc.37.3.347.16044

Chapter 2: Decision Support for Capacitated Vehicle Routing

with Backtracking Restrictions

Decision Support for Capacitated Vehicle Routing with Backtracking Restrictions

Abstract

This paper introduces a model for the Capacitated Vehicle Routing Problem (CVRP) that provides a novel, robust approach to limit the amount of backtracking allowed in the solution. This problem is motivated by the increase in vehicles purchased via online marketplaces such as eBay and an associated increase in the direct consumer procurement of shipping services for transporting recently purchased vehicles. We approach this problem in the context of a profit seeking objective while considering the cognitive processes and behavioral preferences of the driver as important to the ultimate solution of the routing problem. We offer a method for producing a set of good solutions that are differentiated based on backtracking characteristics of the directional flow of the route.

2.0 Introduction

Direct consumer procurement of a shipping provider for large items (such as a car), is a relatively new and rapidly growing phenomenon. Today many consumers arrange to have large items shipped in open, online, and often auction-driven marketplaces. A lucrative market exists for shipping providers to meet the increased demand of shipping these large items. We are therefore motivated to study how shipping providers should select jobs and associated routes to fulfill customers' demands in order to maximize profit. Additionally, we propose the basis for a decision support system (DSS) to allow shipping providers to select from multiple available routes with a novel, robust approach to limit the amount of backtracking allowed in the final solution.

We were inspired to consider this problem by the relatively new phenomenon of online car shopping on websites such as eBay Motors. Thousands of vehicles are sold on eBay Motors every day and an increasing number of these sales are taking place with large distances between the buyer and seller. In such cases, the buyer has two options to receive the vehicle: 1) travel to the vehicle and drive it home, or, 2) have the vehicle shipped to them. We focus on the latter option which has seen a growth in popularity in recent years (uShip.com, 2015).

Five or ten years ago, this type of purchase was fairly rare for the general consumer; however, today it occurs hundreds or even thousands of times per day. eBay Motors alone has facilitated the sale of over 5.2 million vehicles online (eBay Inc., 2015). During the 3rd quarter of 2012, the Gross Merchandise Value (GMV) of vehicle sold on eBay Motors was approximately $2 billion. Of the sales in Q3 of 2012 as well as in Q3 of 2014, 77% were interstate transactions increasing the likelihood of the need to ship the newly purchased vehicle (eBay Inc., 2012, 2015).

Broadly speaking, we attempt to determine the optimal set of arcs in a road network for a delivery vehicle to traverse. As such, this type of problem is a variant of the vehicle routing problem (VRP), which has become a classic optimization problem. The VRP has been studied in the Operations Research literature since its introduction in 1959 (Dantzig & Ramser, 1959). Since 1959, VRP research has both enhanced our ability to model and solve this type of problem as well as broadened the definition of what constitutes a VRP through extensions to the original model. Online shipping procurement offers a new opportunity to broaden the VRP through the unique aspects of the problem that must be considered and modeled.

The decision problem we address involves the profit maximizing selection of the jobs to accept and delivery route to take given the capacity constraints on the delivery vehicle, the pickup and delivery points of the non-fungible items associated with each job, the pickup and delivery time windows associated with each job, and a user-defined backtracking limitation (discussed below). A large array of research has investigated various individual components that make up this problem; however, we have been unable to find published research that addresses the specific combination of issues considered here.

The remainder of our paper is organized as follows. We first review the related literature on the VRP and its many variants. Next, we introduce the issue of backtracking within a VRP and discuss the limitations this phenomenon presents in routing models and solutions that ignore it. We then introduce our new MILP formulation for this problem, which allows the user to define the amount to backtracking that is permissible. We provide an example illustrating the advantages of our model from a decision support standpoint. We then introduce and solve a set of test problems with different parameters and characteristics to investigate the computational performance of the proposed MILP formulation under a variety of data conditions and compare

its performance to a heuristic solution method. Finally, we present conclusions and opportunities for future research in this area.

2.1 Literature Review

As mentioned previously, the VRP is a classic optimization problem and, as such, has volumes of literature describing different formulations, solutions, and aspects of the problem. A review of the VRP, both in terms of recent formulations and solutions is offered by Toth and Vigo (Toth & Vigo, 2002b) as well as Cordeau et al. (2002). A specific look at exact solutions to the VRP under capacity and time constraints is offered by (Baldacci, Mingozzi, & Roberti, 2012).

The problem we address combines elements of several different VRP variants including:

- Capacitated Vehicle Routing Problem (CVRP)

- Vehicle Routing with Pickups and Deliveries (VRPPD)

- Pickup and Delivery Selection Problem (PDSP)

- General Pickup and Delivery Problem (GPDP)

The Capacitated Vehicle Routing Problem (CVRP)

The CVRP is the problem of selecting routes for a fleet of trucks that simultaneously meet the demand of customers while abiding by the capacity of the trucks. Trucks in a CVRP are loaded at a depot location and then travel to the customer sites. Typically, other constraints are also added such as time windows, truck types, and possibly optional deliveries. The typical objective function of the CVRP is to minimize the cost of serving the selected routes. The CVRP is a well-studied problem in the literature that has recently been studied with respect to exact algorithms (Baldacci et al., 2012), demand uncertainty (Gounaris, Wiesemann, & Floudas,

2013), and optimizing with respect to specifically minimizing fuel consumption as opposed to cost or distance (Xiao, Zhao, Kaku, & Xu, 2012). Additional overview references pertaining to the CVRP are available in (Golden, Raghavan, & Wasil, 2008; Toth & Vigo, 2002a).

The Vehicle Routing Problem with Pickups and Deliveries (VRPPD)

In the VRPPD, the goal is to select an optimal set of routes to meet the demand of customers at minimal cost. However in the VRPPD, each transportation request has an origin and a destination rather than the origin always being a depot location. Typically, models assume that all requests must be fulfilled and ignore situations where job selection would be required (as required in this paper). Frequently, loads being shipped in the VRPPD are handled as commodities in that there is no specific destination for an item once it is picked up, it just will need to be delivered to some location that has a need for the item (i.e. a delivery node).

Obviously, in the case of shipping cars, each car has specific pickup and delivery locations. In the VRPPD literature, our problem is most closely related to the Single Vehicle One-to-One Pickup and Delivery Problem (SVPDP) (Cordeau, Laporte, & Ropke, 2008). The VRPPD is a well-studied problem with over 30 years of published research support (Berbeglia, Cordeau, Gribkovskaia, & Laporte, 2007). As such, many surveys and generalized discussions exist including (Savelsbergh & Sol, 1995; Toth & Vigo, 2002b).

The Pickup and Delivery Selection Problem (PDSP)

As the name implies, the PDSP is focused on the optimal selection of transportation requests. In contrast to the CVRP or VRPPD, the typical objective function in the PDSP is to maximize profit by selecting the most profitable jobs. The general focus of the PDSP is on selection of jobs as opposed to routing. However, optimal route selection can be included in the model if such an objective is desired (Maes, Caris, Ramaekers, Janssens, & Bellemans, 2012;

Savelsbergh & Sol, 1995; Schönberger, Kopfer, & Mattfeld, 2003). In the reviewed literature discussing problems modeled as a PDSP or a variant of the PDSP, researchers focus their solution efforts on large problems and heuristic solutions. Therefore, a contribution of this work will be to look at MILP solutions to smaller sized problem.

The General Pickup and Delivery Problem (GPDP)

Savelsbergh and Sol discuss the characteristics of a pickup and delivery problem that differs from a typical VRP (Savelsbergh & Sol, 1995). In addition, they construct a flexible model that can handle the typical instances of pickup and delivery problems and coin the term *General Pickup and Delivery Problem*. While the authors are focused on solving large problems, their model contains many elements needed to model our example problem. In our proposed model, discussed below, we will build upon the GPDP model and add missing components motivated from the three VRP variants discussed above that are required to accurately model our problem of interest.

2.1.1 VRP Decision Support Systems (DSS)

Since the 1980s, researchers have recognized the need to go beyond the mere identification of optimal routes toward an actual DSS for the VRP (Stefanski, Beulens, & Hoefnagels, 1989). The VRP DSS literature ranges from focusing on ease of use and flexibility (Chen, Hwang, Tan, & Lin, 1993) to the incorporation of modern specialized heuristics and constraints (Faiz, Krichen, & Inoubli, 2014; Krichen, Faiz, Tlili, & Tej, 2014). Additionally, the incorporation of geographic information system (GIS) data into a DSS exists both in the research literature (Mendoza, Medaglia, & Velasco, 2009) and in commercial software such as ESRI's ArcGIS.

Our goal is not to implement a complete VRP DSS. Rather we offer a novel approach to obtaining multiple good solutions by controlling the amount of backtracking allowed in a given route. We could find no existing examples of this in our review of the literature.

2.2 Backtracking

In modeling and solving instances of the type of problem described in the introduction, as well as a host of other sample VRPs, we observed that the optimal route often involves reversing course and going in an opposite direction or returning to an area that had been visited before. We refer to this as backtracking and it can be primarily attributed to time windows and capacity constraints that are inherent in this type of problem.

This backtracking phenomenon is often observed as loops in the optimal route. As an example, consider Figures 2-1, 2-2, and 2-3. Figure 2-1 has no backtracking restriction, and includes a large loop in the middle of the route. Figure 2-2 has a moderate backtracking restriction, which eliminates the large loop, but still allows the final route to backtrack toward the origin. Figure 2-3 has a tight backtracking restriction and thus the route flows in a far more direct route from origin to destination.

Figure 2-1: No Backtracking Restriction
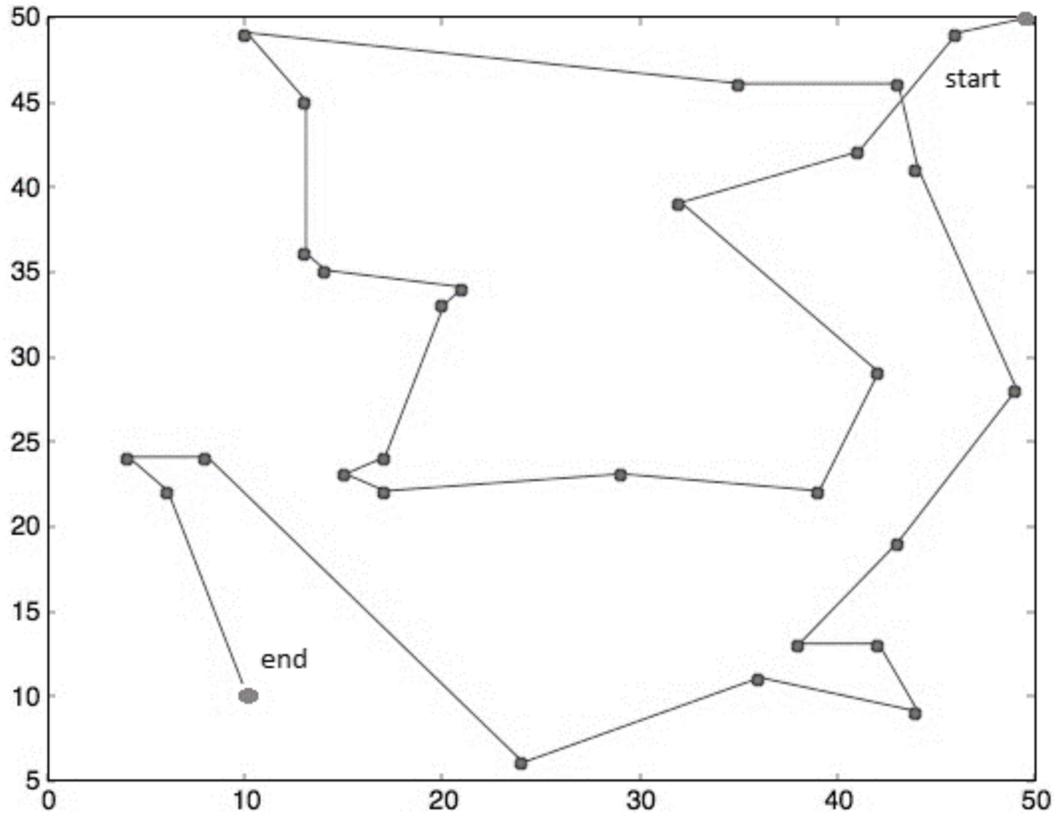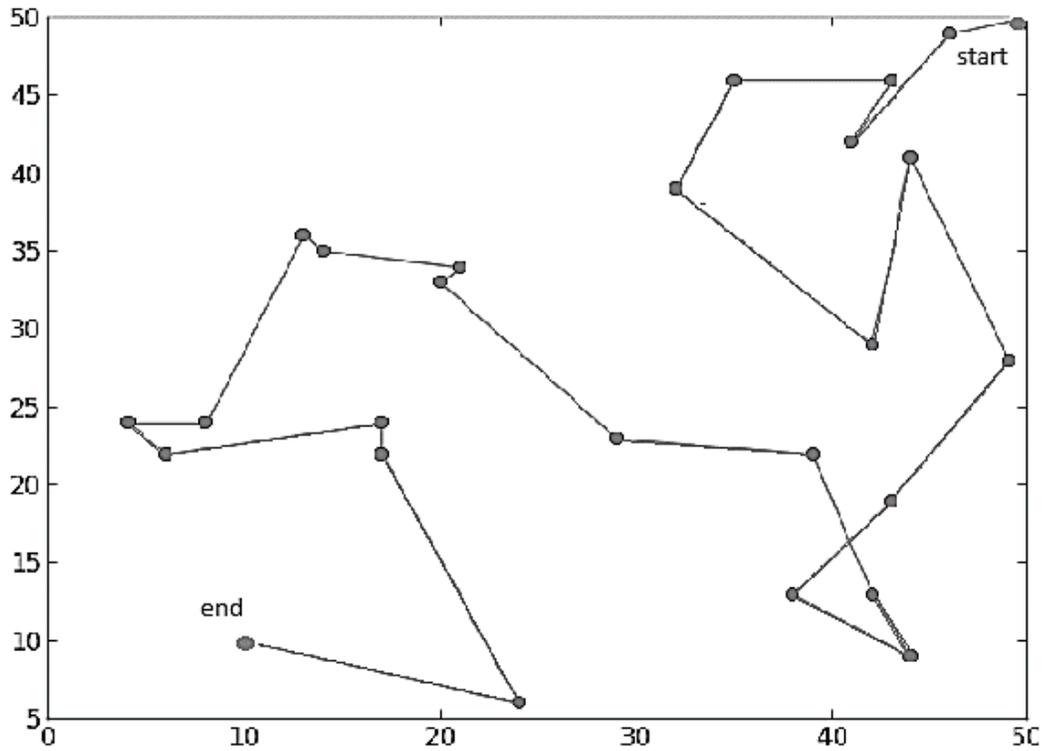
Figure 2-2: Moderate Backtracking Restriction
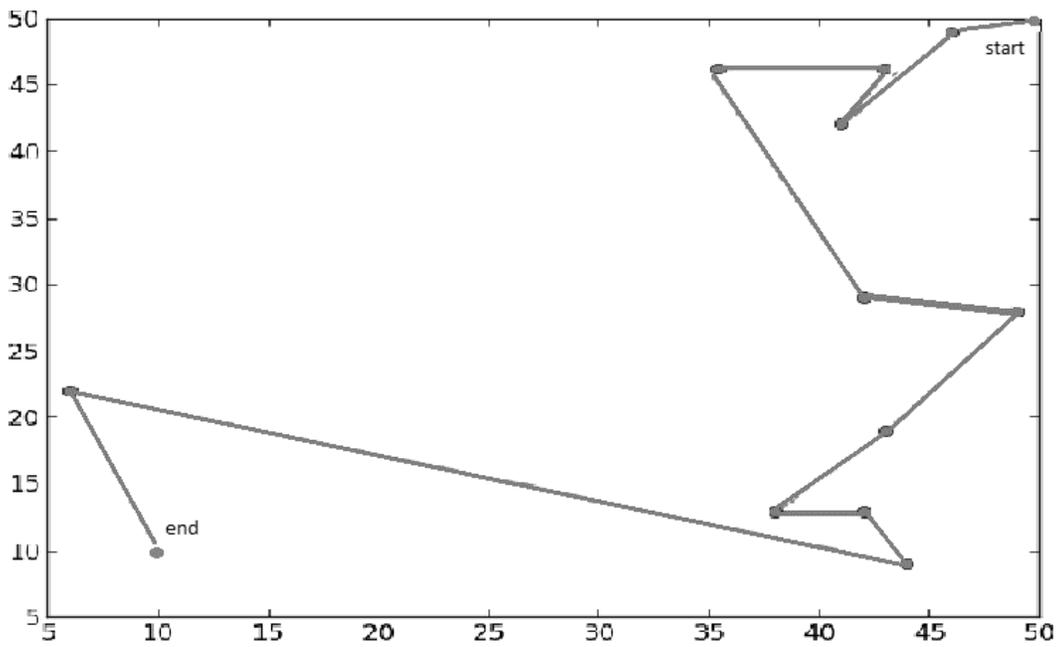
Figure 2-3: Tight Backtracking Restriction

In this example, the tight backtracking restriction illustrated in Figure 3 forces the route to proceed in more of a northeast to southwest flow. Backtracking up (north) or to the right (east) is only allowed if these are relatively short segments in the route.

In our experiments, we defined an allowable distance limit for vertical (north/south) and horizontal (east/west) backtracking. We then used the maximum norm (or L^{∞}-norm) combined with the desired direction of travel to define the boundaries around each network node which would be considered allowable backtracking. Obviously this is only one way of describing backtracking and we devote section 7.0 below to a further discussion on alternatives, extensions, and limitations of this type of approach.

The incorporation of backtracking restrictions in this type of problem is important from a behavioral logistics perspective as these different solutions offer the decision maker another basis for making their decision apart from the objective function or the traditional performance metrics of cost and revenue. For example, suppose a truck driver delivering cars is at the start node in Figure 1 and has a vehicle on board that must be delivered to the end node. That driver likely has a general delivery route in mind, but might be interested opportunities to increase profit by picking up and delivering other cars that require only minor deviations from the intended route. The routes in Figures 2 and 3 might be more desirable to the trucker than the route in Figure 1, even if they are somewhat less profitable.

As another example, if one of the possible deliveries sends the driver into a populous city during rush-hour, they might opt to forgo that job even if the maximum profit would be achieved by taking the job. In the simplest form, this type of preference may generalize into a

parameterized constraint that can be entered into the model where, effectively, the constraint would enforce a condition such as "avoid rush hour traffic."

As a final example, consider the situation where a driver wants to consider other activities that cannot be accomplished until the chosen route is completed (e.g., getting home in time for a child's birthday party). The factors that contribute to deciding whether or not to forgo this other activity may very well be more complex than simply revenue minus cost. This is an example of what Gino and Pisano refer to as a complex system where human behavior is a central driver (2008).

As these examples demonstrate, there are more complex behavioral issues in VRPs that depend on more factors than can readily be factored into a model. As a result, we submit that this problem is ripe for behavioral issues to influence which route truly is the best route for the decision maker, presenting the opportunity for a harvest of new research in this area.

The VRP is often studied from the standpoint of a dispatcher determining routes for a large fleet of trucks. However, existing literature suggests that 90% of shipping providers have a small fleet of less than 6 trucks and over 250,000 truck drivers in the US are self-employed (Sun, 2013). Additionally, a recent survey conducted on drivers in Texas, Ontario, and Indiana found that 65% of truck drivers can freely choose their own routes during the planning stage (Toledo et al., 2013). This motivates us to frame the problem with the view that the driver of the truck is the decision maker; or where the driver, effectively, is also the dispatcher. Thus, we view the cognitive patterns and behavioral preferences of the driver as important to the ultimate solution of the routing problem.

With these issues in mind, we consider how one might reduce the frequency of backtracking loops and develop a set of constraints that limit the distance the optimal solution is

allowed to deviate from the general direction of the route. In addition to reducing backtracking frequency, it also allows us to offer multiple solutions by modifying the distance limit which would aid in the creation of a decision support system (DSS). This DSS will allow us to offer the shipping provider not just one optimal solution but rather a set of good solutions from which to choose.

2.3 Mathematical Model

We model the VRP with backtracking restrictions by modifying the GPDP formulated by Savelsbergh and Sol (1995). In keeping with their notation, we let N be the set of transportation requests (*i.e.* cars we will consider shipping). For each request $i \in N$ we know the origin location N_{i+} as well as the destination N_{i-}. In our model, each pickup location has a load quantity, q_i, of positive one and each destination location will have a load quantity of negative one. We let N_+ and N_- be the set of all pickup and delivery locations respectively and let the set $V = \{N_+ \cup N_-\}$. Let set W contain our starting and ending locations, k^+ and k^-, respectively, for the overall trip. Finally, for all arcs $(i, j) \in V \cup W$ let d_{ij} denote the travel distance, t_{ij} the travel time, and c_{ij} the travel cost.

We have two types of jobs: 1) jobs that ***must*** be fulfilled due to prior commitments, and 2) jobs that we have the ***option*** to fulfill. Therefore, let S_c denote the set of cars for which firm pickup/delivery *commitments* have already been made: $S_c = \{ l \mid$ car l *must* be picked up and delivered$\}$. Let S_p denote the set of *potential* cars we *may* decide to pick up and deliver: $S_p = \{ l \mid$ car l *may* be picked up and delivered$\}$.

In order to control backtracking in our model we introduce a parameter τ which will be the maximum distance in the backwards direction that will be allowed. We modify τ to produce

the moderate and tight backtracking restrictions in the example described in section 3.0 and in the remaining discussion below. We then let B_i be the set of locations that are greater than τ away from node i and then let J_i be the index of the transportation request associated with location i.

We let R be the set of cars on-board when leaving starting location k^+; R= { l | car l is on-board at k^+}. We then let:

- x_{ij} = a binary variable which has a value of 1 if arc (i,j) is in our final route; and 0 otherwise.

- δ_l = a binary variable with a value of 1 if car l is chosen for pickup and delivery; and 0 otherwise.

- $Y_i \geq 0$ specifies the load on the truck when arriving at node i.

- $D_i \geq 0$ specifies the departure time for leaving node i.

We also define the following parameters:

- r_l is the revenue gained for fulfilling the job associated with car l;

- L_i is the close of the time window for node i;

- Q is the capacity of (or maximum number of cars allowed on) the truck;

Our formulation of the problem is then given as follows:

Maximize Profit:
$$\sum_{i \in N} r_i \delta_i - \sum_{(i,j) \in V \cup W} c_{ij} x_{ij}$$

(1)

Subject To:
$$\sum_{j \in V \cup W} x_{mj} = \sum_{j \in V \cup W} x_{jm} = \delta_i, \quad \forall\, i \in N, m \in N_{i+}\, \cup N_{i-}$$

(2)

$$\sum_{j \in V \cup k-} x_{k+,j} = 1 \qquad (3)$$

$$\sum_{j \in V \cup k+} x_{jk-} = 1 \qquad (4)$$

$$D_{k+} = 0 \qquad (5)$$

$$D_P \leq D_q, \qquad \forall\, i \in N, p \in N_{i+}, q \in N_{i-} \qquad (6)$$

$$D_i + t_{ij} \leq D_j + M\left(1 - x_{ij}\right) \forall i,j \in V \cup W \qquad (7)$$

$$Y_{k+} = \sum_{l \in R} \delta_l \qquad (8)$$

$$Y_i \leq Q, \qquad \forall\, i \in V \cup W \qquad (9)$$

$$-M\left(1 - x_{ij,}\right) + Y_j \leq Y_i + q_i \leq Y_j + M(1 - x_{ij}), \qquad (10)$$

$$\forall\, i,j \in V \cup W$$

$$D_i \leq L_i, \qquad \forall\, i \in V \cup W \qquad (11)$$

$$\delta_l = 1, \qquad \forall\, l \in S_c \qquad (12)$$

$$D_i - D_b + M\delta_l \leq M, \qquad \forall\, b \in B_i, \delta_l \in J_i \,\,\forall\, i \in V \cup W \qquad (13)$$

$$X_{ij} \in \{0,1\}, \qquad \forall\, i,j \in V \cup W \qquad (14)$$

$$\delta_l \in \{0,1\}, \qquad \forall\, l \in N \qquad (15)$$

$$D_i \geq 0, \qquad \forall\, i \in V \cup W \qquad (16)$$

$$Y_i \geq 0, \qquad \forall\, i \in V \cup W \qquad (17)$$

The objective in (1) is to maximize profit, which will consist of a revenue component and a cost component. The revenue is simply a given parameter: if we pick up and deliver car l (i.e.

$\delta_l = 1$) we will receive the revenue r_l associated with car l. The cost component in the objective is the estimated cost of traveling between the selected nodes without specifically taking into account the fuel consumption costs associated with a particular load. According to Franzese and Davidson (2011), the addition of a relatively light load of a single car will only make a small difference compared to the cost that will be incurred simply by driving the truck. In other words, trucks are quite efficient when it comes to hauling loads.

Constraint (2) ensures that if the truck enters a node, it also leaves that node and will only enter a node if the associated job is accepted (*i.e.* $\delta_l = 1$). Constraints (3) and (4) ensure that the starting location and ending locations are included in the solution. Constraint (5) ensures that the departure time for leaving the starting node ($k+$) is set to 0. Constraint (6) ensures that the decision variable representing the departure time of the pickup location (D_p) is less than the departure time of the destination location (D_q).

Constraint (7) requires that if we are going from node i to node j (i.e. $X_{ij} = 1$), then the departure time of i, plus the travel time to node j, is less than the departure time of node j. This constraint is a linearization of the following: $X_{ij} = 1 \Rightarrow D_i + t_{ij} \leq D_j$, $\forall i, j \in V \cup W$.

Constraint (8) sets the initial load (Y_{k+}) to the summation of all loads that are currently on board when we leave the starting node ($k+$). Constraint (9) ensures that the decision variable representing the total load when arriving at a particular node (Y_i), is less than the vehicle capacity (parameter Q).

Constraint (10) ensures that we accurately represent the total load when arriving at a particular node (Y_j) includes the load change (q_i) from the previous node. Recall that q_i can be either positive (representing a pickup) or negative (representing a delivery). Since load values at

two nodes are only applicable if a route between those two nodes is actually used (X_{ij} =1),

constraint (10) is a linearization of the following: $X_{ij} = 1 \Rightarrow Y_i + q_i = Y_j, \ \forall \ i,j \in V \ \cup \ W.$

Constraint (11) enforces the time windows, requiring each departure time to be prior to

the time window closing. Constraint (12) ensures that the binary variable associated with a

committed job is set to 1.

Constraint (13) eliminates backtracking by utilizing the departure time of a given location

that is already inherent in our model. For instance, if we have to backtrack to go from location j

to i, we ensure that the departure time of location i (D_i) is less than that of j (D_j) but only if the

job associated with location j has been chosen to be accepted. Finally, constraints (14) – (16)

define our variables that were mentioned in the discussion above.

2.4 Example Problem Solution

We implemented this model as an MILP based on the following scenario. Suppose we are

driving from New York, NY to Mobile, AL with 5 cars initially on board with a truck capacity is

6 vehicles. We observed actual vehicle shipments from uShip.com to generate the potential jobs,

revenue, and time window parameters. We calculated the distance between each city in our

model (all $(i,j) \in V \ \cup \ W$) and used that distance as the basis for the cost components in the

objective function. Our objective then is to select the set of jobs and the associated route that will

maximize profit per the job candidates listed in Table 2-1.

Car	From	To	Revenue ($)	Time Window (pickup)	Time Window (deliver)
1	-onboard-	Raleigh, North Carolina	3120	-	9011
2	-onboard-	Winston–Salem, North Carolina	3426	-	100000
3	-onboard-	Chesapeake, Virginia	2226	-	100000
4	-onboard-	Baltimore, Maryland	1146	-	1368

5	-onboard-	Virginia Beach, Virginia	2184	-	9024
6	Charlotte, North Carolina	New Orleans, Louisiana	800	7500	8858
7	Memphis, Tennessee	Baton Rouge, Louisiana	740	5332	100000
8	Washington D.C.	Nashville, Tennessee	975	5048	8823
9	Norfolk, Virginia	Augusta, Georgia	850	100000	100000
10	Greensboro, North Carolina	Little Rock, Arkansas	800	100000	100000
11	Pittsburgh, Pennsylvania	Durham, North Carolina	1150	8813	100000
12	Louisville, Kentucky	Lexington, Kentucky	1450	7139	7273
13	Richmond, Virginia	Montgomery, Alabama	725	100000	100000
14	Philadelphia, Pennsylvania	Fayetteville, North Carolina	450	6573	100000
15	Atlanta, Georgia	Mobile, Alabama	550	8366	100000

Table 2-1: Car Shipments

With no restrictions on backtracking, the optimal solution produces a profit of $9849 by fulfilling all jobs except 11 and 13 following the route in Figure 2-4 below.
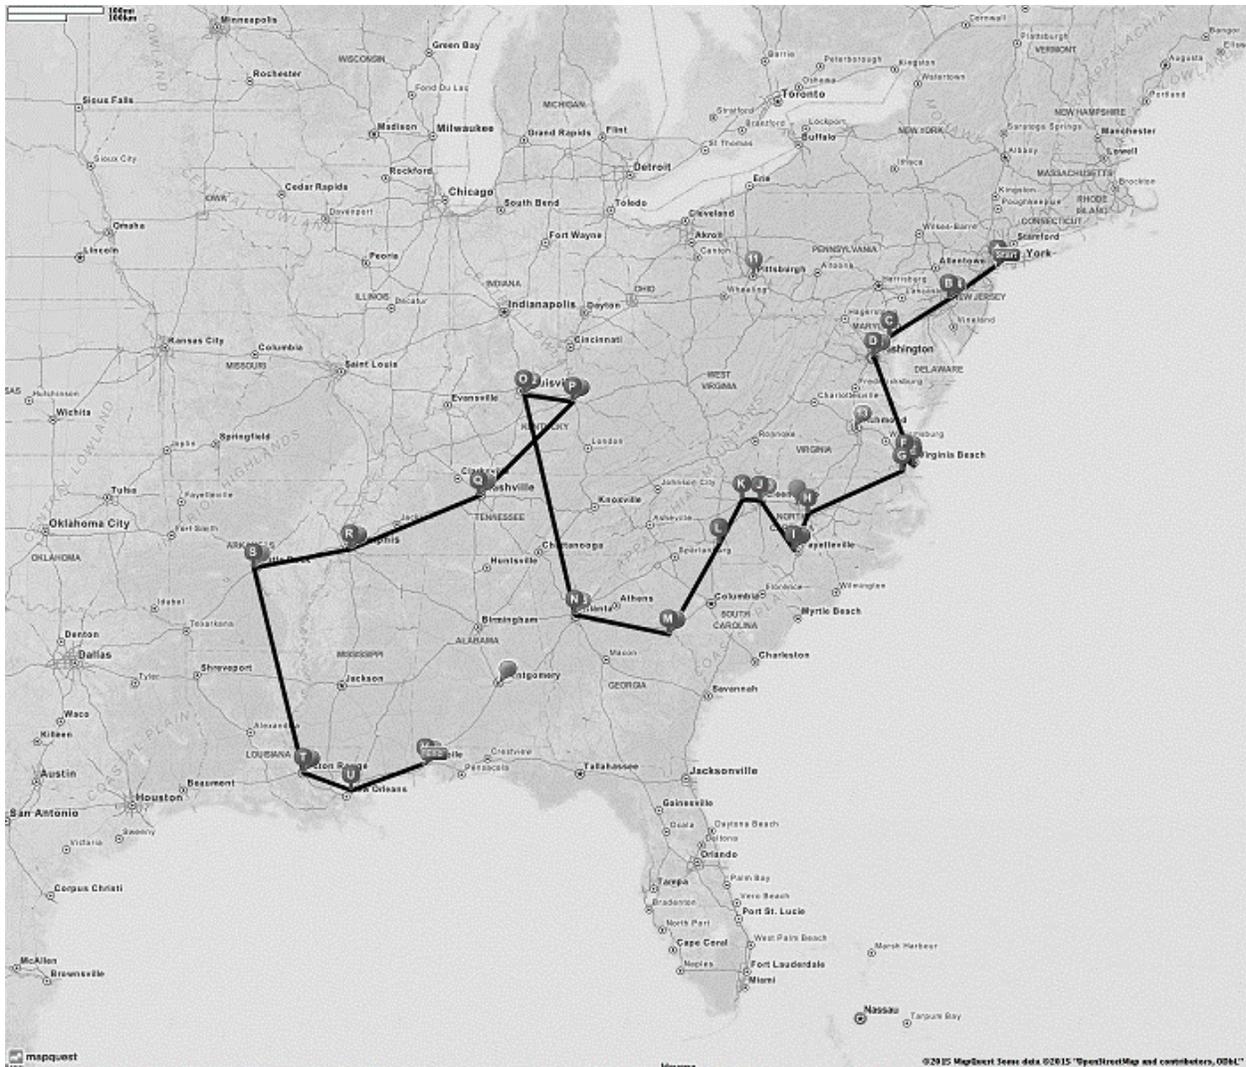
Figure 2-4: No Backtracking Restrictions

The segment from Atlanta to Louisville is an example of a route that we consider to be backtracking. The general flow of the route is southwestern (e.g., New York to Alabama) but this segment is an obvious change to the north. Clearly, this segment is worth it from a pure profit point of view, otherwise it would not have been included in the optimal solution. However, it could very well be that this route travels many more miles in order to gain just a few dollars extra in profit.

One potential strategy to eliminate backtracking is to add constraints that disallows certain segments; for instance, Atlanta to Louisville. However, if one adds that constraint, the

optimal route includes Atlanta to Nashville to Louisville. So we still ended up with backtracking, it is just more gradual.

In testing the performance of the strategy presented in section 1.3 above, changing the amount of backtracking allowed, we were able to successfully influence the backtracking and introduce a new performance metric, profit per mile, which could be used by a decision maker to compare the different routing options. The results are summarized in Table 2-2 below.

Backtracking Restriction	Profit	Miles	Profit per Mile	Jobs fulfilled
None	9849	2956	3.331867	13
Moderate	9794	2491	3.931754	12
Tight	9688	1663	5.825616	9

Table 2-2: Backtracking Comparison Results

The results show that even though no backtracking restriction produced the most profit, it was not a tremendous amount of profit above the solutions with moderate and tight backtracking restrictions. Figures 2-5 and 2-6 below show the routes taken with the moderate and tight backtracking restrictions respectively. Incorporating these three route options, along with the performance metrics of profit, miles, and the derived profit per mile would allow the decision maker to make an informed comparison and decide whether the extra miles is worth the added profit. In the sections that follow, we conduct a further analysis on the effect of backtracking restrictions on the solution obtained by our model.
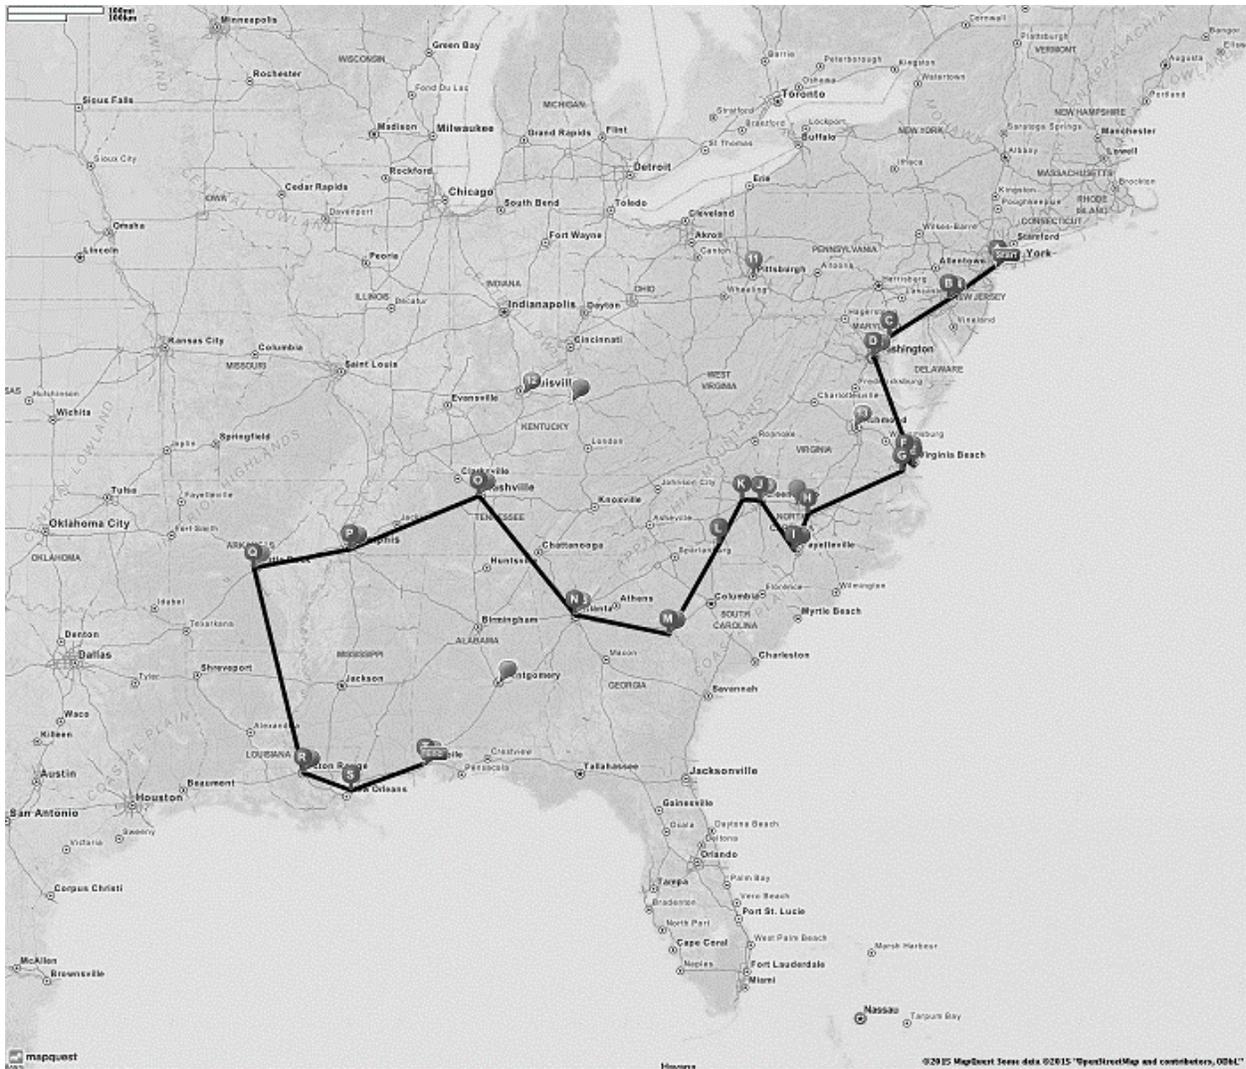
Figure 2-5: Moderate Backtracking Restriction
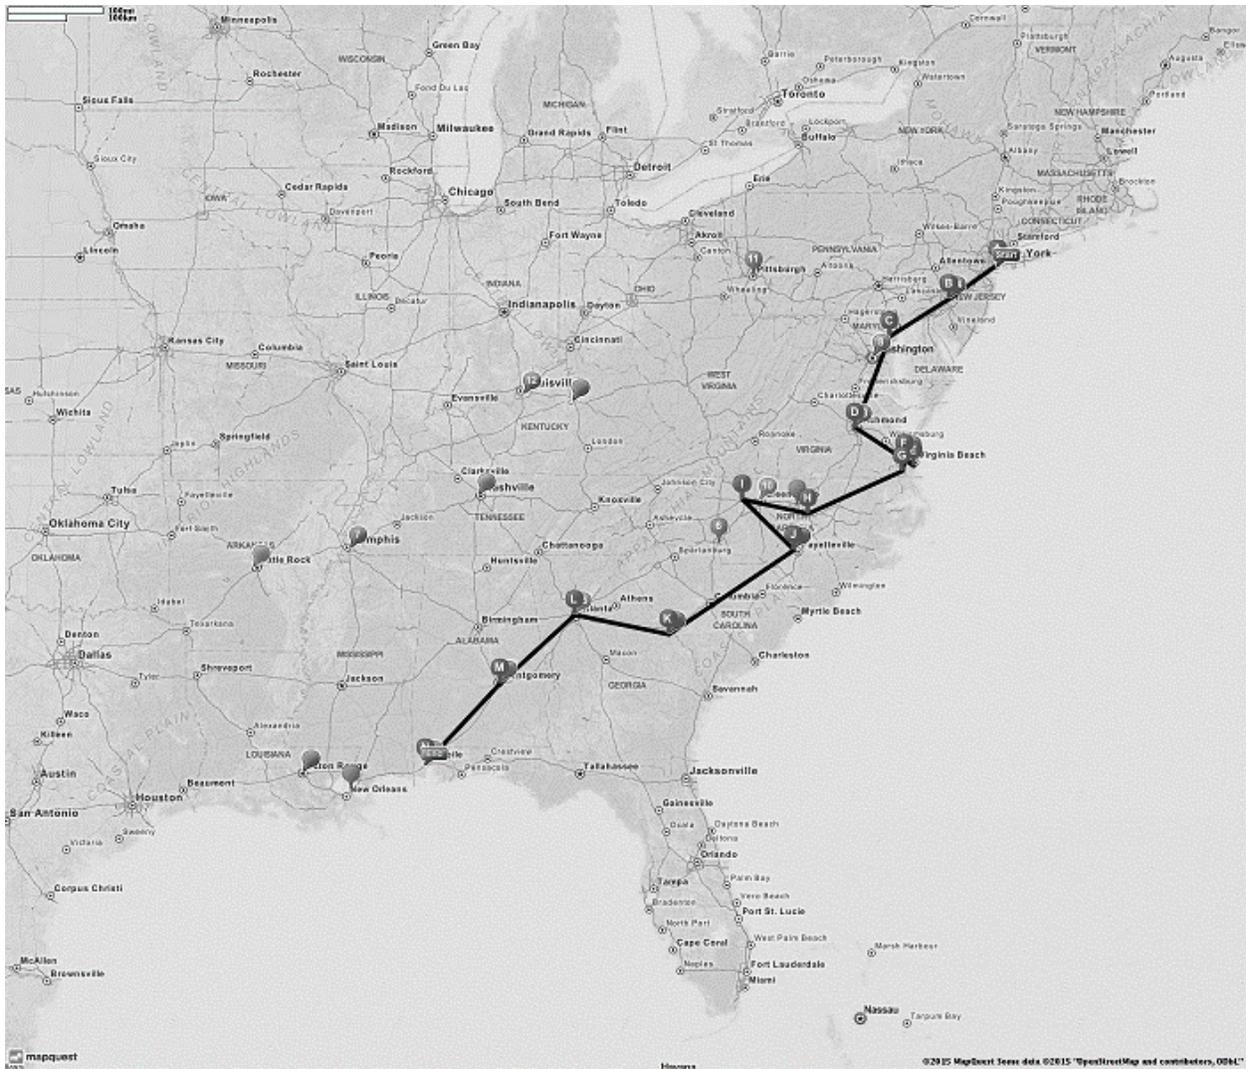
Figure 2-6: Tight Backtracking Restriction

2.5 Computational Analysis

To the best of our knowledge, there exist no published benchmarks for our problem to which we can compare results. We were motivated by two sets of benchmark problems. The first was created by Ropke and Cordeau (2009) in their study of a new algorithm to solve the Pickup and Delivery Problem with Time Windows. The second set was created by Salhi and Nagy (1999) whose problems were modifications of a classic VRP benchmark dataset published in

Christofides, Mingozzi, and Toth (1979). We used these two sets of benchmark problems as the starting point for developing our own set of benchmark problems specific to our problem of interest.

We needed to develop benchmark problems that contained the following information: location of cities, paired cities associated with a particular car, the time windows associated with each pickup and delivery, the revenue associated with each car, and whether or not the car was on-board at the beginning of the problem. The combination of these made utilizing existing benchmarks quite difficult, particularly when considering how they can be dependent upon each other. For example, the pairing of the pickup and delivery cities for a car will have an impact on the time windows for each (i.e. the delivery city likely has a time window that closes after the pickup city). Additionally since we are specifically interested in influencing the amount of backtracking in a given route, we needed problems that had the same directionality to them. In our case, we chose to start north and east and end south and west. This is not a limitation of our model, rather it was chosen for comparability and convenience. All of this motivated us to use mainly the city location information from the benchmark problems discussed above and create the city pairing, time windows, and decide which cars were already on board.

Using this strategy, we created 30 problems that we solved for our analysis. Since we are primarily interested in studying the backtracking component of the problem, we fixed the following parameters:

- On-board jobs at start: 1

- Capacity of the truck: 10

- Jobs to consider: 15; implying 30 cities total

The 30 instances described above were solved using IBM ILOG CPLEX version 12.6.0.0 running on a 3.5 GHz 6-core Intel Xeon E5 processor, running Apple OS X 10.9.4. We experimented with different CPLEX options and concluded that changing from the default values resulted in no definitive improvement to either the solutions or the solution times. Therefore, with the exception of the maximum runtime, which we configured for 600 seconds, we used the default CPLEX options. We solved each problem four times, changing only the backtracking restriction and the results are averaged and presented in Table 2-3 below.

Backtracking (τ)	Avg Profit	Avg Miles	Avg Profit per mile	Avg Gap (%)
15 mi	$770.7	130.99	5.91	0.02
30 mi	$1,268.2	210.98	6.12	8.72
45 mi	$1,429.3	220.49	6.62	4.19
No limit	$1,453.6	227.17	6.57	4.01

Table 2-3: Backtracking Experiment Results

These results do not paint the same picture as the example problem presented previously. The profit differential between the tightest backtracking restriction and no restriction is nearly 100%. However when we compare the 30 and 45 mile backtracking restrictions with no restriction, we see the profit differential reduced significantly. The box and whiskers plot shown in Figure 2-7 below illustrates how small the difference is.
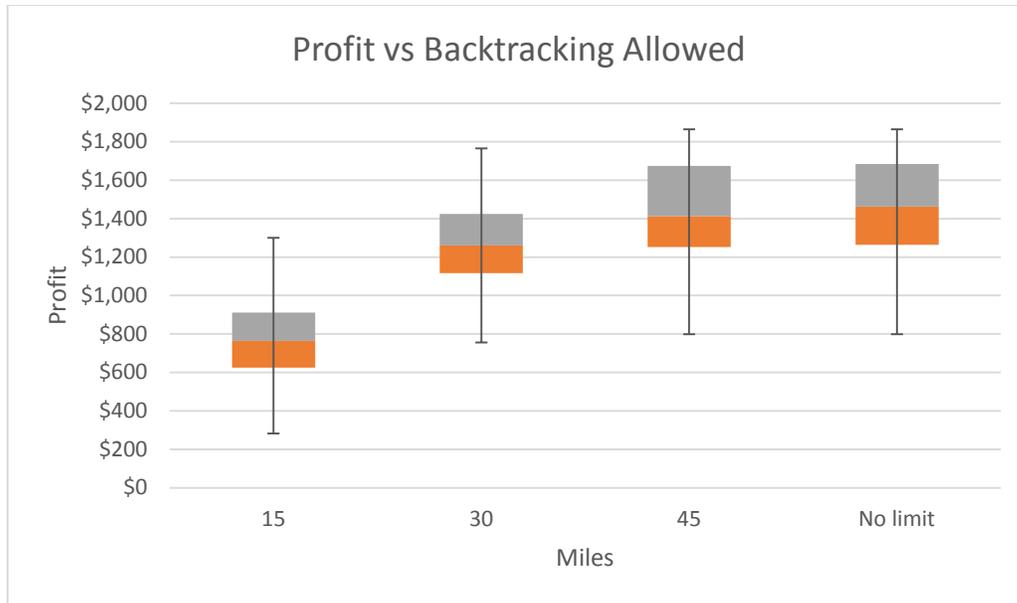
Figure 2-7: Summary of Tradeoff Between Profit and Backtracking

As we consider the performance indicator of profit per mile, we also see a narrow range across all 4 backtracking parameter settings. This implies that a decision maker could benefit by seeing at least the 30, 45, and unlimited backtracking restriction scenarios to choose from.

2.6 Flexibility of Backtracking Restrictions

In this paper, we have used a rather simplistic but effective backtracking restriction by looking at the allowable distance and direction along the horizontal and vertical axes individually. In other words, we used direction and the maximum norm to define the backtracking boundary around each location i. In the example problem, this type of backtracking restriction made sense and was relatively easy to implement. There likely exist other problems and domains where different or perhaps more complex backtracking restrictions would be desirable.

Since we are using an L^p-norm (where $p=\infty$ thus far), we can support other p-norms simply by varying the value of p to adjust the contour around the location i. The L^1-norm for instance would allow us to use the straight-line distance between points as the backtracking restriction. Figure 2-8 shows some of the common contours that can be created by modifying p. Up to this point, we have assumed there was a desired direction of travel that was important and have thus considered that aspect of the relationship; however, there is no requirement for that to be the case.

1a: Unit sphere for ℓ_1 norm   1b: Unit sphere for ℓ_2 norm   1c: Unit sphere for ℓ_∞ norm   1d: Unit sphere for ℓ_p quasinorm
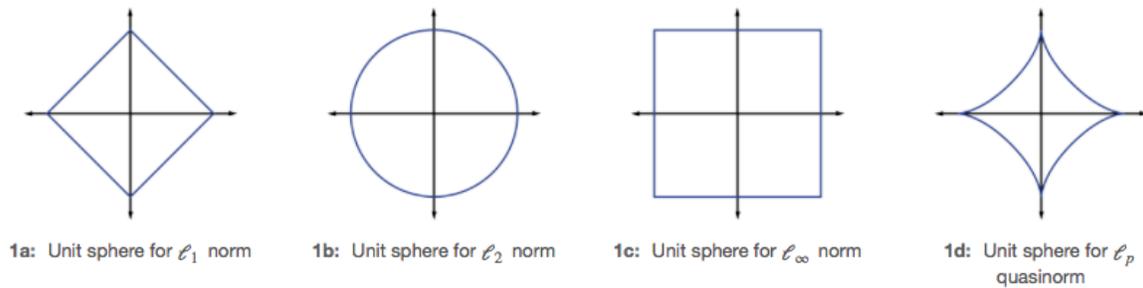
Figure 2-8: Contours From Different p Norms (Baraniuk, Davenport, Duarte, & Hedge, 2011)

Our model could be easily modified to support node specific backtracking restrictions by allowing for different values of τ at each location i (τ_i). It is also possible to support different backtracking restriction functions at each location. The principle requirement is that for each location i we must be able to identify the set of locations (B_i) that would exceed the allowed backtracking limit.

2.6.1 Multi-Directional Backtracking Restrictions

A more difficult problem to address is that of supporting more than one desired direction of travel. For example, consider the problem of designing a route that begins and ends at the same node. Breaking such a problem into two independent (out and back) parts is easily shown to be suboptimal. Another approach to this challenge is to mirror the problem so that the flow

appears to be unidirectional. Figure 2-9 shows an example route with two directions of flow; essentially a circular route. In the backtracking restrictions we have discussed above, traveling between nodes 5 and 6 would likely be considered backtracking.
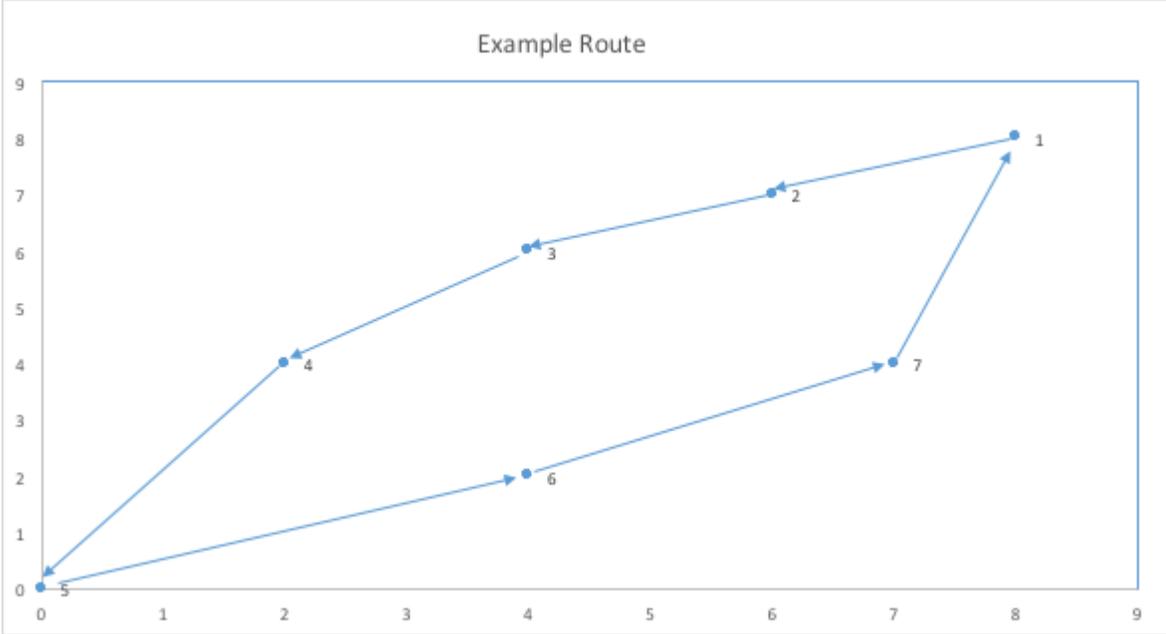
Figure 2-9: Circular Route

Figure 2-10 illustrates how the same route may be transformed so that the general flow stays in one direction. The weakness to this type of approach is that it requires doubling the problem size and will likely be much more difficult to solve. This transformation may lend itself well to a heuristic type of solution (e.g. genetic algorithm, tabu search, etc.). In any event, techniques for dealing with different types of backtracking situations offers a new and interesting area for additional research in this domain.
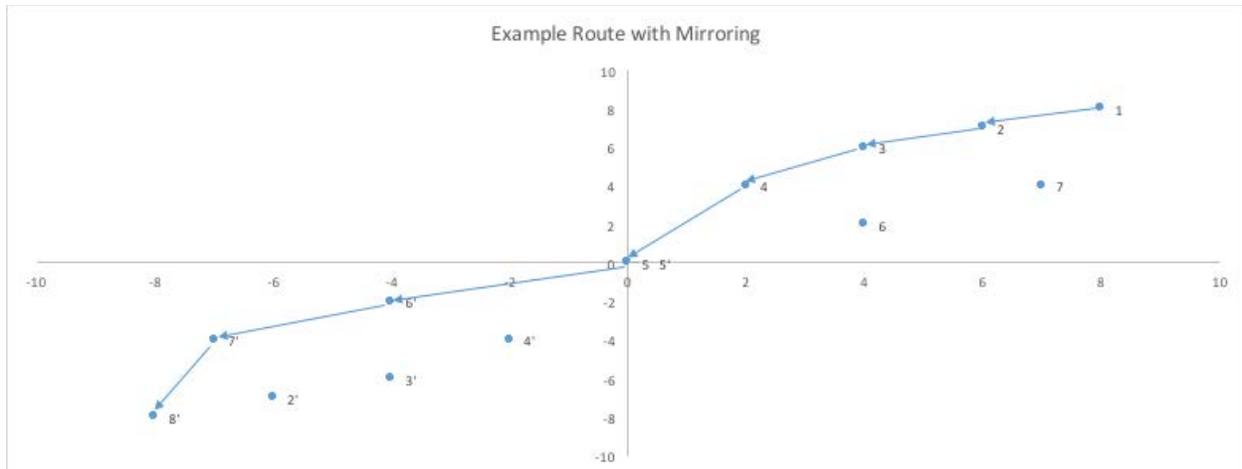
Figure 2-10: Mirrored Route

2.7 Conclusions and Future Research

We investigated a car shipping problem with backtracking restrictions by introducing the problem, modifying existing MILP formulations to model it, and then solved 30 newly introduced benchmark problems. These problems were intended to represent problems that an independent trucker may encounter and thus were relatively small when compared to less constrained VRPs. However, solving smaller problems does not equate to providing little to no value. Not only does studying smaller problems lead to insights in larger, more complex problems (Cordeau et al., 2008), but as reported above, 90% of shipping providers have a small fleet of less than 6 trucks and over 250,000 truck drivers in the US are self-employed (Sun, 2013). This indicates that drivers would likely be interested in improving their route selection and perhaps more importantly be pleased to have options to choose from. This combined with the growth in shipping cars purchased via eBay leads us to conclude that these types of problems should be further studied.

Our problem formulation and solution methodology has no dependencies on large computational resources. An independent trucker could use the same techniques and find a good route and job selection on a laptop computer. A larger system, fronted by a web interface could solve many similar problems for a host of independent truckers using the same formulation and backed by cloud computing resources.

One of the challenges to solving a real-world problem similar to the one we study in this paper is to limit the number of jobs considered, which directly affects the problem size. A reliable way to filter jobs that have little to no chance of being included in the solution should be developed. It is possible that other strategies for implementing the backtracking restriction could also help to reduce the problem size and thus other backtracking restriction techniques should be explored.

The problem could also consider the social and environmental costs of emissions in the final solution. Existing work in the literature lays the groundwork for this type of consideration within our model (Park, Rakha, Farzaneh, Zietsman, & Lee, 2010).

Finally, the multiple solutions produced by modifying the backtracking restriction parameter would serve a VRP DSS nicely. The development of or integration into an existing DSS should be explored.

References

Baldacci, R., Mingozzi, A., & Roberti, R. (2012). Recent exact algorithms for solving the vehicle routing problem under capacity and time window constraints. *European Journal of Operational Research*, *218*(1), 1–6. http://doi.org/10.1016/j.ejor.2011.07.037

Baraniuk, R., Davenport, M., Duarte, M., & Hedge, C. (2011). *An Introduction to Compressive Sensing*. OpenStax CNX. Retrieved from http://cnx.org/contents/f70b6ba0-b9f0-460f-8828-e8fc6179e65f@5.7

Berbeglia, G., Cordeau, J.-F., Gribkovskaia, I., & Laporte, G. (2007). Static pickup and delivery problems: a classification scheme and survey. *TOP*, *15*(1), 1–31. http://doi.org/10.1007/s11750-007-0009-0

Chen, A.-P., Hwang, C.-H., Tan, G. L. H., & Lin, C.-Y. (1993). Development of a decision support system for service delivery. *Computational Economics*, *6*(2), 115–129. http://doi.org/10.1007/BF01299230

Christofides, N., Mingozzi, A., & Toth, P. (1979). The vehicle routing problem. In N. Christofides, A. Mingozzi, P. Toth, & C. Sandi (Eds.), *Combinatorial Optimization* (pp. 315–338). Chichester, UK: Wiley.

Cordeau, J.-F., Gendreau, M., Laporte, G., Potvin, J.-Y., & Semet, F. (2002). A Guide to Vehicle Routing Heuristics. *The Journal of the Operational Research Society*, *53*(5), 512–522. http://doi.org/10.2307/823019

Cordeau, J.-F., Laporte, G., & Ropke, S. (2008). Recent Models and Algorithms for One-to-One Pickup and Delivery Problems. In B. Golden, S. Raghavan, & E. Wasil (Eds.), *The Vehicle Routing Problem: Latest Advances and New Challenges* (pp. 327–357). Springer US. Retrieved from http://link.springer.com/chapter/10.1007/978-0-387-77778-8_15

Dantzig, G. B., & Ramser, J. H. (1959). The Truck Dispatching Problem. *Management Science*, *6*(1), 80–91. http://doi.org/10.1287/mnsc.6.1.80

eBay Inc. 2013. "eBay Motors Fast Facts Q3 2012." Accessed September 25. http://legacy.ebayinc.com/assets/pdf/fact_sheet/eBay_Motors_Fact_Sheet_Q32012.pdf.

eBay Inc. 2015. "eBay Marketplaces Fast Facts At-A-Glance (Q3 2014)." Accessed February 6. http://www.ebayinc.com/sites/default/files/MP_FastFacts_Q32014_0.pdf.

Faiz, S., Krichen, S., & Inoubli, W. (2014). A DSS based on GIS and Tabu search for solving the CVRP: The Tunisian case. *The Egyptian Journal of Remote Sensing and Space Science*, *17*(1), 105–110. http://doi.org/10.1016/j.ejrs.2013.10.001

Franzese, O., & Davidson, D. (2011). Effect of Weight and Roadway Grade on the Fuel Economy of Class-8 Freight Trucks. *Tech Rep ORNL/TM-2011/471, Oak Ridge National Laboratory*. Retrieved from http://info.ornl.gov/sites/publications/files/Pub33386.pdf

Gino, F., & Pisano, G. (2008). Toward a Theory of Behavioral Operations. *Manufacturing & Service Operations Management*, *10*(4), 676–691. http://doi.org/10.1287/msom.1070.0205

Golden, B. L., Raghavan, S., & Wasil, E. A. (2008). *The Vehicle Routing Problem: Latest Advances and New Challenges*. Springer.

Gounaris, C. E., Wiesemann, W., & Floudas, C. A. (2013). The Robust Capacitated Vehicle Routing Problem Under Demand Uncertainty. *Operations Research*. http://doi.org/10.1287/opre.1120.1136

Krichen, S., Faiz, S., Tlili, T., & Tej, K. (2014). Tabu-based GIS for solving the vehicle routing problem. *Expert Systems with Applications*, *41*(14), 6483–6493. http://doi.org/10.1016/j.eswa.2014.03.028

Maes, T., Caris, A., Ramaekers, K., Janssens, G. K., & Bellemans, T. (2012). Pickup and delivery selection problem. Retrieved from https://uhdspace.uhasselt.be/dspace/handle/1942/14365

Mendoza, J. E., Medaglia, A. L., & Velasco, N. (2009). An evolutionary-based decision support system for vehicle routing: The case of a public utility. *Decision Support Systems*, *46*(3), 730–742. http://doi.org/10.1016/j.dss.2008.11.019

Park, S., Rakha, H., Farzaneh, M., Zietsman, J., & Lee, D.-W. (2010). Development of fuel and emission models for high speed heavy duty trucks, light duty trucks, and light duty vehicles. In *2010 13th International IEEE Conference on Intelligent Transportation Systems (ITSC)* (pp. 25–32). http://doi.org/10.1109/ITSC.2010.5624972

Ropke, S., & Cordeau, J.-F. (2009). Branch and Cut and Price for the Pickup and Delivery Problem with Time Windows. *Transportation Science*, *43*(3), 267–286. http://doi.org/10.1287/trsc.1090.0272

Salhi, S., & Nagy, G. (1999). A Cluster Insertion Heuristic for Single and Multiple Depot Vehicle Routing Problems with Backhauling. *The Journal of the Operational Research Society*, *50*(10), 1034–1042. http://doi.org/10.2307/3009928

Savelsbergh, M. W. P., & Sol, M. (1995). The General Pickup and Delivery Problem. *Transportation Science*, *29*(1), 17–29. http://doi.org/10.1287/trsc.29.1.17

Schönberger, J., Kopfer, H., & Mattfeld, D. C. (2003). A Combined Approach to Solve the Pickup and Delivery Selection Problem. In P. D. U. Leopold-Wildburger, P. D. F. Rendl, & P. D. G. Wäscher (Eds.), *Operations Research Proceedings 2002* (pp. 150–155). Springer Berlin Heidelberg. Retrieved from http://link.springer.com/chapter/10.1007/978-3-642-55537-4_24

Stefanski, M. M., Beulens, A. J. M., & Hoefnagels, J. P. (1989). From Vehicle Routing Packages Towards Routing DSS. In P. D. D. Pressmar, D. K. E. Jäger, P. D. H. Krallmann, P. D. H. Schellhaas, & P. D. L. Streitferdt (Eds.), *Operations Research Proceedings 1988* (pp.

571–578). Springer Berlin Heidelberg. Retrieved from http://link.springer.com/chapter/10.1007/978-3-642-74862-2_150

Sun, Y. (2013). *Decision making process and factors routing* (Thesis). Massachusetts Institute of Technology. Retrieved from http://dspace.mit.edu/handle/1721.1/82853

Toledo, T., Sun, Y., Rosa, K., Ben-Akiva, M., Flanagan, K., Sanchez, R., & Spissu, E. (2013). Decision-Making Process and Factors Affecting Truck Routing. In M. B. Akiva, H. Meersman, & E. V. D. Voorde, *Freight Transport Modelling* (pp. 233–248). Emerald Group Publishing.

Toth, Paolo, and Daniele Vigo. 2002b. "Models, Relaxations and Exact Approaches for the Capacitated Vehicle Routing Problem." *Discrete Applied Mathematics* 123 (1–3): 487–512. doi:10.1016/S0166-218X(01)00351-1.

Toth, P., & Vigo, D. (2002b). *The Vehicle Routing Problem*. SIAM.

uShip.com. (2015, September 1). *uShip.com - About*. Retrieved from http://www.uship.com/about/

Xiao, Y., Zhao, Q., Kaku, I., & Xu, Y. (2012). Development of a fuel consumption optimization model for the capacitated vehicle routing problem. *Computers & Operations Research*, *39*(7), 1419–1431. http://doi.org/10.1016/j.cor.2011.08.013

Chapter 3: The Truck Driver Scheduling Problem with Fatigue Monitoring

The Truck Driver Scheduling Problem with Fatigue Monitoring

Abstract

Studies by the National Highway Traffic Safety Administration (NHTSA) show that 100,000 police-reported crashes, 71,000 injuries, and 1,550 deaths occur due to drowsy driving each year in the United States (Rau, 2005). However, no model currently exists that incorporates a measure of drowsiness or fatigue into the Truck Driver Scheduling Problem (TDSP). We introduce the first fatigue-aware model for determining the optimal schedule for a driver while maintaining an acceptable level of alertness as well as abiding by time windows and hours of service (HOS) regulations.

3.0 Introduction

Driver fatigue has been empirically identified as a major factor in vehicle crashes (Chen & Xie, 2014). Studies by the National Highway Traffic Safety Administration (NHTSA) show that 100,000 police-reported crashes, 71,000 injuries, and 1,550 deaths occur due to drowsy driving each year in the United States (Rau 2005). Bowman et al. (2012) conclude that driver fatigue is the probable cause of 30% of crashes. According to the National Sleep Foundation, these statistics "significantly underestimate the problem" (National Sleep Foundation 2007) and there exist naturalistic driving studies to support the claim of underestimation (Dingus et al. 2006; Klauer et al. 2006).

For commercial trucks, many countries (and the European Union) have hours of service (HOS) regulations that attempt to control the amount of fatigue a truck driver experiences while driving (Cappuccio, Miller, & Lockley, 2010, pp. 424–429; "Summary of Hours of Service Regulations," 2013). When these regulations are modified, studies are published attempting to quantify the effect of the changes on safety outcomes (Blanco et al., 2011; Hanowski, Hickman, Olson, & Bocanegra, 2009). These regulations attempt to balance the safety benefits with the added costs that will be incurred by compliance. Quantifying these benefits, costs, and other factors introduced with regulatory compliance is an area of active research. However, the currently available results point to significant opportunity for additional research, particularly with respect enforcing regulations that produce the desired safety benefits without being overly burdensome (Heaton, 2005; Jensen & Dahl, 2009; Kemp, Kopp, & Kemp, 2013).

Mathematical models and solution techniques for the Truck Driver Scheduling Problem (TDSP) and Vehicle Routing Problem (VRP) incorporating HOS constraints have been

introduced in the literature (Archetti & Savelsbergh, 2009; Goel, 2012a, 2014; Xu, Chen, Rajagopal, & Arunapuram, 2003). Goel and Vidal (2014) model and compare several countries' HOS constraints from both a cost and risk standpoint. However, we have been unable to find a model that incorporates fatigue into the model itself.

Mathematical models for predicting alertness have existed at least since Borbely's two-process model was published in 1982 (Borbely, 1982). Since then several additional models have been published, many of which are derived from Borbely's initial work. See Mallis et al. (2004), Gundel et al. (2007), and Dawson et al. (2011) for a review of these models. The accuracy of these models has been shown to be similar (Van Dongen, 2004). Specifically with respect to vehicle crashes, it has been shown that one can reasonably predict crashes with a mathematical model that uses a sleep/wake predictor that is based on the Three Process Model of Alertness (Åkerstedt, Connor, Gray, & Kecklund, 2008). The Three Process Model of Alertness (TPMA) is itself an extension of Borbely's original model.

We introduce the Truck Driver Scheduling Problem with Fatigue Monitoring (TDSPFM) by incorporating the TPMA. This would allow us to identify schedules where the driver is likely to remain alert and thereby reduce the likelihood of being involved in a crash. Additionally, this model will give a comparative basis for measuring the effectiveness of models that incorporate HOS regulations.

3.1 Literature Review

As mentioned above, there exists an established body of research around both the VRP and TDSP with a wide array of variations and additional constraints. Additionally, the literature on fatigue prediction is reasonably well established. What is missing, and the void this paper

seeks to fill, is the merger of these two areas of research to create a mathematical model that can be used to determine optimal driving schedules while maintaining an acceptable level of predicted alertness.

The fact that there is not yet a TDSPFM does not imply that there is no research into reducing driver fatigue. Fatigue detection is an active area of research and product development (Brown, Johnson, & Milavetz, 2013; Chang & Chen, 2014; Luo, Hu, & Fan, 2013; Zhang et al., 2014). In all of the articles we reviewed, the focus was on detecting driver fatigue in the vehicle. Generally speaking, the goal of that research/work is to alert the driver before an accident occurs by using data captured during the trip. For instance, one could use yawn and blinking frequency to predict fatigue using computer vision technologies (Jin, Park, & Lee, 2007) and then alert the driver when the fatigue estimate crosses some pre-defined threshold.

Workplace fatigue has been well-studied and results in "an unsafe condition in the workplace" (Lerman et al., 2012). Entire systems of accounting for and managing the risk that fatigue introduces, aptly named Fatigue Risk Management Systems (FRMS), have been defined and their use has been advocated (Fourie, Holmes, Bourgeois-Bougrine, Hilditch, & Jackson, 2010). FRMS's are outside the scope of this paper, however it is worth noting that our proposed TDSPFM would fit into the existing FRMS framework.

It has also been recognized that the workplace fatigue research would be directly applicable to vehicle crashes (Lerman et al., 2012). The Fatigue Avoidance Scheduling Tool (FAST) based on the Sleep, Activity, Fatigue, and Task Effectiveness model has been developed for the aviation and rail transportation industries, but also has applicability to driver scheduling (Hursh, Balkin, Miller, & Eddy, 2004). The North American Fatigue Management Program

provides training to address issues of driver fatigue by way of online courses and presentations (North American Fatigue Management Program, 2012).

3.1.1 Alertness Models

The TPMA model utilized in this research consists of three primary processes that have been previously published (Åkerstedt, Folkard, & Portin, 2004) and are described briefly below. Process C represents the circadian influence on alertness; this process encompasses the effect that the time of day can have on sleepiness. Process S describes the exponential decline in alertness as a function of the time awake. This decline is then reversed in process S' which describes recovery as a function of the time asleep. Figure 3-1 below shows the way the above processes effect alertness at different times during the day. Process W describes the lack of alertness at the time of waking up. Since we will assume driving does not take place immediately following waking up, we ignore the W component of the TPMA model for this research. Finally, since the original publication of the TPMA model, another process has been added (Åkerstedt, Axelsson, & Kecklund, 2007). This process is U, which stands for "ultradian" and explains an afternoon dip in alertness.
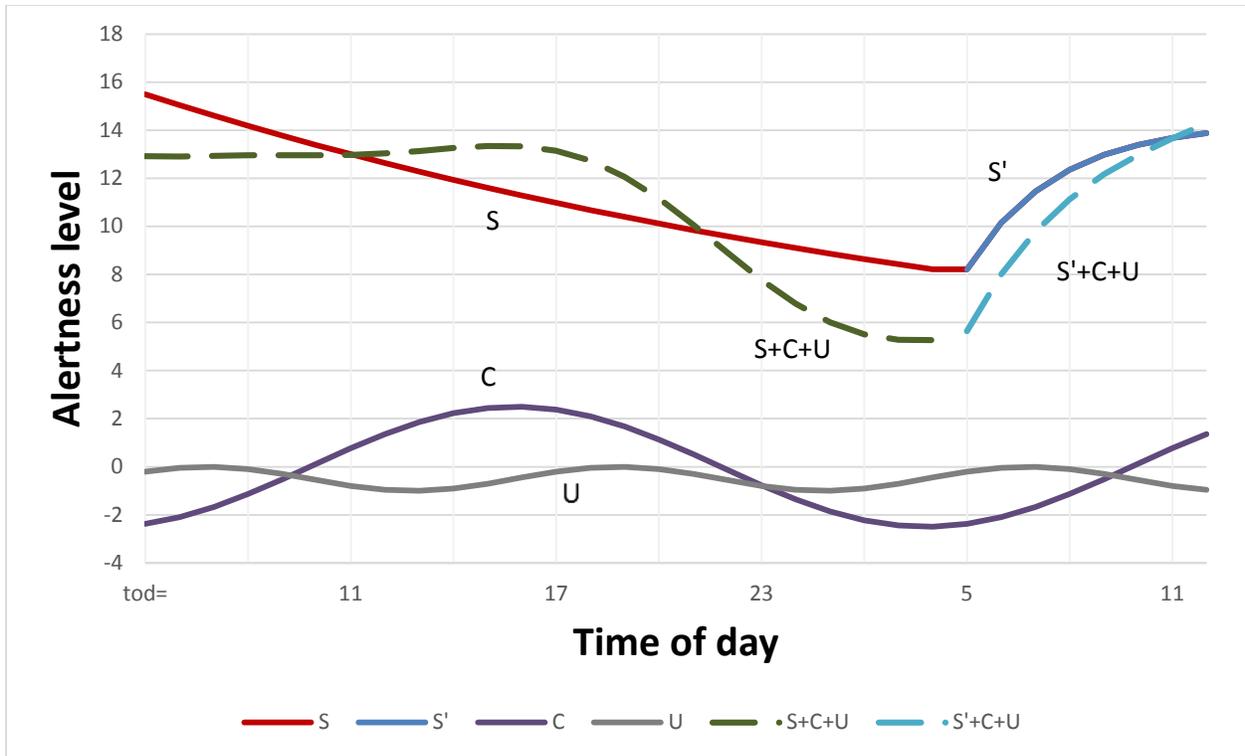
Figure 3-1: Components of the TPMA

For the implementation of our TDSPFM model, we will use the validated TPMA model and parameters presented in (Ingre et al., 2014). Therefore, our model predicts alertness while driving as the summation of S+C+U. This produces an alertness score with values ranging from 1 to 21. According to the sleep literature, a TPMA value (or score) of "3" corresponds to extreme sleepiness, while "14" represents high alertness, and "7" to a borderline sleepiness threshold (Åkerstedt et al., 2004). While a minimum allowed alertness score is a required parameter in our model, the appropriate value of this parameter is a matter for further research.

3.1.2 Sleep Assumptions

When implementing the TDSPFM in this paper, we make some conservative estimates regarding sleep. First, we assume that the driver is well-rested and alert when they start the work week. Second, we do not consider caffeine or other drug use that could affect how sleepy or alert a driver is. Additionally, we assume that when the driver takes a long rest break, they get an uninterrupted period of recovering sleep. In other words, we do not factor in things like sleep disorders, noise and other distractions inhibiting sleep, or drivers that choose to do things other than sleep during the time when they could be sleeping.

It is worth noting that the TPMA allows one to predict sleep schedules and to factor in those results into the alertness recovery process (S'). For this paper, we will not take advantage of this functionality in favor of simplifying the driver scheduling process. However, it is an area to consider for future research and the reader is encouraged to see (Åkerstedt et al., 2008) for more information.

3.1.3 Hours of Service Regulations

In general, HOS regulations function primarily by imposing rules related to how long drivers may stay on the road, the conditions relating to rest frequency, and the duration/types of rest periods. As a result, a driver's schedule for a particular day might look like Figure 3-2 below.

Figure 3-2: A Truck Driver Schedule (modified from Goel, 2012b)

For the purposes of the remainder of this paper, we consider the HOS regulations in the United States. Specifically we use the parameters defined and explained in Goel (2014). Although we use the US regulations, it is worth noting that other HOS restrictions could easily be accommodated by our model.

3.2 Mathematical Model

We now present our model for the TDSPFM with a goal of finding the schedule that has the minimal duration while maintaining an acceptable level of alertness, abiding by HOS regulations, and complying with the time windows at each location on the route. As in (Goel, 2012b) and (Goel, 2012a), we assume that all rest breaks occur at stops along the route. Since the sequence of locations and driving time between locations is fixed, the primary decisions to be made are those regarding the duration of rest periods (r_i) at each location.

We consider a sequence of N locations which will be visited by a truck driver. Each location $i \in N$ has a time window and some duration of work associated with it. We let the opening time window of each location be zero for the sake of simplicity and focus on the closing time window of L_i for each $i \in N$. The work duration is denoted as w_i for each location $i \in N$. We let $t_{i,i+1}$ be the amount of time it takes to drive between locations i and the next location. The arrival time and departure time of each location i is denoted as A_i and D_i.

57

Since the sequence of locations in our model is fixed, our focus is on choosing the optimal set of rest times at each location in order to minimize duration while maintaining a minimum level of alertness. Rest times are denoted below as r_i for each location $i \in N$. The duration we wish to minimize is calculated as the difference between the arrival time at the final location and the departure time at the first location.

As mentioned earlier, we chose to use parameters representative of the HOS regulations in the United States. These parameters and their values used for our implementation of the TPMA are displayed in Table 3-1 below and come from Ingre et al. (2014).

Parameter Name	Value (h)	Description
t^{rest}	10	Minimum rest time to be considered a long break
t^{break}	.5	Minimum Break
$t^{horizon}$	168	Planning Horizon (1 week)
t^{drive}	11	Maximum drive time since last break
$t^{elapsed|R}$	14	Maximum time since last long rest
$t^{elapsed|B}$	8	Maximum time since last break
$t^{maxsleep}$	10	Maximum amount of sleep allowed
$t^{awakedelay}$	0.5	The delay after waking before driving can begin
$t^{sleepdelay}$	1	The delay after rest begins before falling asleep
t^{max_rest}	16	Maximum rest period allowed
la	2.4	Lower asymptote of the internal alertness scale
d	-0.0353	Decay in alertness

g	-0.38135645	Recovery multiplier
ha	14.3	Higher asymptote of the internal alertness scale
bl	12.5	Break level of recovery function S'
Cm	0	Mesor of process C
Ca	2.5	Amplitude of process C
p	16.8	Default circadian phase
Um	0.5	Mesor of process U
Ua	0.5	Amplitude of process U

Table 3-1: Model Parameters

In order to abide by the HOS regulations, we need to keep track of the following variables which are presented along with their associated equations:

- Time driving upon arriving at location i, dt_i:

$$dt_i = \begin{cases} dt_{i-1} + t_{i-1,i}, & r_{i-1} < t^{rest} \\ t_{i-1,i} & , & r_{i-1} \geq t^{rest} \end{cases} , \forall\, i \in N$$

- Time working since last long rest upon arriving at location i, st_i:

$$st_i = \begin{cases} st_{i-1} + t_{i-1,i} + w_i, & r_{i-1} < t^{rest} \\ t_{i-1,i} + w_i & , & r_{i-1} \geq t^{rest} \end{cases} , \forall\, i \in N$$

- Time since last break upon arriving at location i, stb_i:

$$stb_i = \begin{cases} stb_{i-1} + t_{i-1,i} + w_i, & r_{i-1} < t^{break} \\ t_{i-1,i} + w_i & , & r_{i-1} \geq t^{break} \end{cases} , \forall\, i \in N$$

To determine the alertness of the driver at a given location i ($alertness_i$), we need to introduce the following variables:

- The time awake upon arriving at location i, taw_i:

$$taw_i = \begin{cases} taw_{i-1} + t_{i-1,i} + r_{i-1} + w_i & , \quad r_{i-1} < t^{rest} \\ t_{i-1,i} + (r_{i-1} - t^{maxsleep}) + t^{awakedelay} + w_i, & r_{i-1} \geq t^{rest} \end{cases} , \forall i \in N$$

- The time of day upon arriving at location i, tod_i:

$$tod_i = A_i \bmod 24 \ \forall i \in N$$

Finally, we show the calculations related to the components of the TPMA which will ultimately allow us to compute an alertness score at a given location i ($alertness_i$):

- Process S at location i, S_i:

$$S_i = la + (S'_{i-1} - la) * e^{d*taw_i} \ \forall i \in N$$

- Subprocess SB_i, used to determine the proper S'_i to use:

$$SB_i = ha - (ha - ss)e^{g(t^{maxsleep} - t^{sleepdelay})} \ \forall i \in N$$

- Process S' at location i, S'_i:

$$S'_i = \begin{cases} ha - (ha - bl)e^{g(t^{maxsleep} - t^{sleepdelay} - bt_i)} & , \quad r_{i-1} \geq t^{rest} \ AND \ SB_i \geq bl \\ S_i + g(t^{maxsleep} - t^{sleepdelay}) * (bl - ha), & r_{i-1} \geq t^{rest} \ AND \ SB_i < bl \ , \forall i \in N \\ S'_i & , \quad r_{i-1} < t^{rest} \end{cases}$$

- Process C at location i, C_i:

$$C_i = Cm + Ca * \cos\left(\left(2\frac{\pi}{24}\right) * (tod_i - p)\right) \ \forall i \in N$$

- Process U at location i, U_i:

$$U_i = Um + Ua * \cos\left(\left(2\frac{\pi}{12}\right) * (tod_i - p - 3)\right) \quad \forall i \in N$$

The above processes allow us to compute the alertness score at a given location i (*alertness$_i$*):

$$alertness_i = S_i + C_i + U_i, \quad \forall i \in N$$

Recall from the TPMA discussion above, the process components that make up the alertness score are non-linear. Therefore, it is very likely that the minimum alertness score will not happen at a given location i, but rather during the drive between *i-1* and *i*. We use the pseudo-code presented in Figure 3-3 below to calculate the minimum alertness score (*minalertness$_i$)* along a given route segment.

// step through times from leaving the previous location (i-1) to
//   arriving at current location (i) in 15 minute increments
for t_{now} in D_{i-1} to A_i step 0.25:
    $tod_{now} = t_{now}$ mod 24 // time of day for C and U calculations
    $taw_{now} = taw_{i-1} + t_{now}$ // time awake

    $S_{now} = la + (S'_{i-1} - la) * e^{d*taw_{now}}$

    $C_{now} = Cm + Ca * \cos\left(\left(2\frac{\pi}{24}\right) * (tod_{now} - p)\right)$

    $U_{now} = Um + Ua * \cos\left(\left(2\frac{\pi}{12}\right) * (tod_{now} - p - 3)\right)$

    $alertness_{now} = S_{now} + C_{now} + U_{now}$

    // keep track of minimum alertness
    if $alertness_{now} < minalertness_i$ then
        $minalertness_i = alertness_{now}$

Figure 3-3: Pseudo-code of Minimum Alertness Calculation

Our formulation of the TDSPFM is then given as follows,

Minimize:

$$A_{last} - D_{first} \qquad (1)$$

Subject To:

$$A_i + r_i + w_i = D_i \quad \forall i \in N \qquad (2)$$

$$D_i + t_{i,i+1} = A_{i+1} \quad \forall i \in N \qquad (3)$$

$$dt_i \leq t^{drive}, \forall i \in N \qquad (4)$$

$$st_i \leq t^{elapsed|R}, \forall i \in N \qquad (5)$$

$$stb_i \leq t^{elapsed|B}, \forall i \in N \qquad (6)$$

$$A_{last} + W_{last} \leq t^{horizon} \qquad (7)$$

$$0 \leq r_i \leq t^{max\_rest} \quad \forall i \in N \qquad (8)$$

$$A_i \leq L_i, \quad \forall i \in N \qquad (9)$$

$$minalertness_i \geq TPMA^{min}, \quad \forall i \in N \qquad (10)$$

$$r_i \geq 0, \quad \forall i \in N \qquad (11)$$

The objective in (1) is to minimize the duration of our route from our first location until we arrive at the last location. Constraints (2) and (3) ensure our arrival and departure times capture the time consumed at each location. Constraints (4)-(7) enforce HOS regulations.

Constraint (4) is the HOS drive time constraint. Constraint (5) is the HOS working time constraint. Constraint (6) ensures that a short rest break is taken in accordance with HOS regulations. Constraint (7) ensures that the arrival time and associated work time of the last location (A_{last} and w_{last}, respectively) occur within the time horizon specified for the problem. The parameter $t^{\text{max\_rest}}$ specifies the maximum rest time and constraint (8) ensures that the rest time at each location (r_i) is within that limit.

Constraint (9) is the time window restriction ensuring that the arrival time of each location is prior to the time window of that location (L_i) closing. For the sake of simplicity, we allowed an opening time window at each location to be 0 though an additional constraint to account for an opening time window can easily be added.

Constraint (10) allows us to specify an alertness threshold ($TPMA^{\text{min}}$) and ensures that the alertness score stays above that value. Should we only want to use the hours of service constraints, we set $TPMA^{\text{min}} = 0$. Finally, constraint (11) ensures the rest time at location i (r_i) is non-negative.

We populate the initial arrival time (A_{first}) to correspond with a 6 a.m. start for the work week. We set a moderate initial alertness score of 10.32 based on the Karolinska Sleepiness Scale (KSS) (Kaida et al., 2006; Shahid, Wilkinson, Marcu, & Shapiro, 2011) and the conversion from KSS to TPMA scores (Ingre et al., 2014). We also set the initial rest time and work time (r_{first}, w_{first} respectively) to be 0 and thus we can calculate the departure time (D_{first}). For the purposes of the results presented below, we will use a work time at all other locations (w_i) of 30 minutes.

Similar to the model in Goel (2012b), the TDSPFM only allows the driver to rest after arrival at a location and completing any work required at the location. Should we want to consider rest areas or other stops designed for non-work related activities, dummy locations along the route can be added to the model with a work time of zero.

Constraint (10) above makes the model non-linear and thus a mathematical program (such as an MILP) is impractical. In order to solve the problem, a heuristic was required and we chose to implement a genetic algorithm (GA) solution. We therefore coded this model in Microsoft Excel 2013 and solved it using the Evolutionary solution method of the built-in Solver. For a stopping condition, we set a maximum run time of 600 seconds. The TDSPFM is considered to be a planning problem, thus we are more interested in finding good solution as opposed to fast run-times. These problems were solved on a computer with 128 GB of RAM, 2 2.60 GHz Intel Xeon processors, running the 64-bit Windows 7 operating system.

Because we already know the travel times and sequence of locations to be visited, we coded a repair function that would force a rest of at least t^{rest} in length if continuing would violate one of the HOS restrictions. As is common in GA formulations, we used penalty functions to capture violations of time windows, HOS regulations, and alertness constraints.

3.3 Results

We created a set of randomly generated benchmark problems in order to compare the TDSPFM's performance over a range of minimum alertness levels. Following the methodology of Goel (2012b), we focus solely on minimizing the duration of the schedule while abiding by the constraints including time windows and HOS restrictions. This admittedly leads to a greedy solution that could be considered overly aggressive when compared to realistic truck driving

schedules, but provides a baseline for comparison. We first solved the problem with no alertness minimum ($TPMA^{\min} = 0$), which results in minimizing the duration of the schedule while abiding by the HOS restrictions described above. These results are identified as "Baseline HOS (0)" in Table 3-2.

Next, we compared the baseline solution with those obtained by setting minimum alertness levels ($TPMA^{\min}$) at 7.07, 8.15, and 9.24. These minimum alertness levels were chosen in accordance with the Karolinska Sleepiness Scale (KSS) levels discussed in Ingre et al. (2014), Shahid et al. (2011), and Kaida et al. (2006). These alertness levels are closer to being considered "sleepy" than being "alert". We assume that it would be unrealistic to have a minimum alertness level set so high that the driver can never be tired at any time. Therefore, we study levels that could be described as:

- tired (as opposed to sleepy), alertness level 7.07
- semi-tired, alertness level 8.15
- not tired, alertness level 9.24

Because we are adding constraints and tightening the solution space, we would anticipate the optimal objective value (of the schedule duration) to increase as the minimum alertness level is increased. A key contribution of this work is that it facilitates the quantification and understanding of the magnitude of this tradeoff.

Table 3-2 below shows the averaged results of duration, minimum alertness, and average alertness for all 30 benchmark problems as well as the worst case minimum alertness observed. In addition to summarizing the results, we conducted ANOVA testing to determine if there were differences between our problem configurations. Once the ANOVA results showed there was a

statistically significant difference, we utilized Tukey-Kramer HSD test ($\alpha = 0.05$) to find where the differences were, comparing the results of different alertness level constraints to the baseline (HOS only constraints).

Problem (alertness level constraint)	Duration (hours)	Minimum Alertness	Worst Case Minimum Alertness	Average Alertness	Duration % Increase Over Baseline	Minimum Alertness % Increase Over Baseline
Baseline HOS (0)	99.20	7.9	7.0	9.9	-	-
HOS (tired)	99.22	7.9	7.1	9.9	0.02%	0.05%
HOS (semi-tired)	100.37	8.3[a]	8.2	10.0	1.18%	5.13%
HOS (not tired)	104.65[a]	9.3[a]	9.2	10.6[a]	5.49%	17.94%

Table 3-2: Averaged Results

a: Indicates statistically significant differences from the Baseline HOS (0) at the 0.05 level.

It is somewhat comforting to see that with just the HOS constraints, the average minimum alertness stays above the TPMA sleepiness threshold of 7, though the worst case scenario was at the 7.0 level. Therefore, it is not surprising to see non-statistically significant increases in either duration or alertness when enforcing an alertness threshold of tired. However, a semi-tired threshold, we see a 5.13% increase in the minimum alertness score with a mere 1.18% increase in duration. The difference in minimum alertness was statistically significant ($p < 0.0001$) however the increase in duration was not statistically significant at the semi-tired threshold. Finally, when the threshold is set to prevent the driver from getting below the "not tired" stage, we observed statistically significant increases in both duration and minimum alertness ($p < 0.0001$).

In this study, average alertness measures ranged from 9.9 to 10.6 which are likely reasonable with regard to safety. They also show a small spread which likely indicates that there are few points along the route where sleepiness reaches critical thresholds.

These results illustrate that the existing HOS regulations do allow for drivers to obey the regulations and remain above the general sleepiness level. Recall that an alertness level of 7 is generally considered the "sleepy" cutoff in the sleep literature. However, that cutoff does not take into account what activity the human subject may be performing at the time. Thus, if the general sleepiness level is 7 then perhaps the acceptable alertness level for someone driving an 80,000 pound vehicle at interstate speeds should be somewhat greater.

3.4 Conclusions and Future Work

An advantage of using TPMA scores in the TDSPF model is that they can be transformed into a prediction of the subjective Karolinska Sleepiness Scale (KSS) which has been shown to be a valid means for measuring sleepiness (Kaida et al., 2006). The KSS is a scale that "measures the subjective level of sleepiness at a particular time during the day" (Shahid et al., 2011) and was introduced by Åkerstedt and Gillberg (1990). Since the KSS can be obtained simply by asking a driver to rate their sleepiness on a 9 point Likert scale, future research could readily validate the schedules proposed by our model as well as the alertness parameter values we use.

As mentioned in the literature review, detecting fatigue in near real-time is an active area of research and product development. We do not dispute the importance of that type of solution towards ensuring the safety of our roadways. However, the trucking industry should also want to do as much as it can to produce schedules that reduce the likelihood that real-time fatigue detectors would be needed. An effective schedule development strategy to ensure both cost

effective scheduling and safety is one that takes fatigue and alertness into account during the planning stage. Forcing a driver to stop and take an unplanned rest due to a fatigue detection by a real-time system will be much more costly than avoiding the fatigue by developing a safer schedule.

This paper introduces the first scheduling model that includes the necessary variables and constraints to enforce a minimum alertness level. The results presented in this paper show that it is possible to produce schedules that meet existing HOS constraints while also meeting a minimum alertness level. Our model allows for an estimate of the magnitude of the increased duration of a route that is incurred in order to increase the alertness level. Thus, this work serves as a starting point for developing more cost effective ways to develop safer schedules for commercial truck drivers and safer highways for the traveling public.

References

Åkerstedt, T., Axelsson, J., & Kecklund, G. (2007). Individual validation of model predictions of sleepiness and sleep hours. *Somnologie - Schlafforschung Und Schlafmedizin*, *11*(3), 169–174. http://doi.org/10.1007/s11818-007-0315-7

Åkerstedt, T., Connor, J., Gray, A., & Kecklund, G. (2008). Predicting road crashes from a mathematical model of alertness regulation—The Sleep/Wake Predictor. *Accident Analysis & Prevention*, *40*(4), 1480–1485. http://doi.org/10.1016/j.aap.2008.03.016

Åkerstedt, T., Folkard, S., & Portin, C. (2004). Predictions from the Three-Process Model of Alertness. *Aviation, Space, and Environmental Medicine*, *75*(3), A75–A83.

Åkerstedt, T., & Gillberg, M. (1990). Subjective and objective sleepiness in the active individual. *The International Journal of Neuroscience*, *52*(1-2), 29–37.

Archetti, C., & Savelsbergh, M. (2009). The Trip Scheduling Problem. *Transportation Science*, *43*(4), 417–431. http://doi.org/10.1287/trsc.1090.0278

Blanco, M., Hanowski, R. J., Olson, R. L., Morgan, J. F., Soccolich, S. A., Wu, S.-C., & Guo, F. (2011). The Impact of Driving, Non-Driving Work, and Rest Breaks on Driving Performance in Commercial Motor Vehicle Operations. Retrieved from http://trid.trb.org/view.aspx?id=1108816

Borbely, A. A. (1982). A Two Process Model of Sleep Regulation. *Human Neurobiol.*, *1*, 195–204.

Bowman, D. S., Schaudt, W. A., & Hanowski, D. R. J. (2012). Advances in Drowsy Driver Assistance Systems Through Data Fusion. In A. Eskandarian (Ed.), *Handbook of Intelligent Vehicles* (pp. 895–912). Springer London. Retrieved from http://link.springer.com/referenceworkentry/10.1007/978-0-85729-085-4_34

Brown, T., Johnson, R., & Milavetz, G. (2013). Identifying Periods of Drowsy Driving Using EEG. *Annals of Advances in Automotive Medicine*, *57*, 99–108.

Cappuccio, F., Miller, D. M. A., & Lockley, S. W. (2010). *Sleep, Health, and Society: From Aetiology to Public Health*. Oxford University Press.

Chang, T.-H., & Chen, Y.-R. (2014). Driver fatigue surveillance via eye detection. In *2014 IEEE 17th International Conference on Intelligent Transportation Systems (ITSC)* (pp. 366–371). http://doi.org/10.1109/ITSC.2014.6957718

Chen, C., & Xie, Y. (2014). The impacts of multiple rest-break periods on commercial truck driver's crash risk. *Journal of Safety Research*, *48*, 87–93. http://doi.org/10.1016/j.jsr.2013.12.003

Dawson, D., Ian Noy, Y., Härmä, M., Åkerstedt, T., & Belenky, G. (2011). Modelling fatigue and the use of fatigue models in work settings. *Accident Analysis & Prevention*, *43*(2), 549–564. http://doi.org/10.1016/j.aap.2009.12.030

Fourie, C., Holmes, A., Bourgeois-Bougrine, S., Hilditch, C., & Jackson, P. (2010). *Fatigue risk Management Systems: A review of the literature*. Citeseer.

Goel, A. (2012a). The Canadian minimum duration truck driver scheduling problem. *Computers & Operations Research*, *39*(10), 2359–2367. http://doi.org/10.1016/j.cor.2011.12.016

Goel, A. (2012b). The minimum duration truck driver scheduling problem. *EURO Journal on Transportation and Logistics*, *1*(4), 285–306. http://doi.org/10.1007/s13676-012-0014-9

Goel, A. (2014). Hours of service regulations in the United States and the 2013 rule change. *Transport Policy*, *33*, 48–55. http://doi.org/10.1016/j.tranpol.2014.02.005

Goel, A., & Vidal, T. (2014). Hours of Service Regulations in Road Freight Transport: An Optimization-Based International Assessment. *Transportation Science*, *48*(3), 391–412. http://doi.org/10.1287/trsc.2013.0477

Gundel, A., Marsalek, K., & Thoren, C. ten. (2007). A critical review of existing mathematical models for alertness. *Somnologie - Schlafforschung Und Schlafmedizin*, *11*(3), 148–156. http://doi.org/10.1007/s11818-007-0312-x

Hanowski, R. J., Hickman, J. S., Olson, R. L., & Bocanegra, J. (2009). Evaluating the 2003 revised hours-of-service regulations for truck drivers: The impact of time-on-task on critical incident risk. *Accident Analysis & Prevention*, *41*(2), 268–275. http://doi.org/10.1016/j.aap.2008.11.007

Heaton, K. (2005). Truck Driver Hours of Service Regulations: The Collision of Policy and Public Health. *Policy, Politics, & Nursing Practice*, *6*(4), 277–284. http://doi.org/10.1177/1527154405282841

Hursh, S. R., Balkin, T. J., Miller, J. C., & Eddy, D. R. (2004). *The Fatigue Avoidance Scheduling Tool: Modeling to Minimize the Effects of Fatigue on Cognitive Performance* (SAE Technical Paper No. 2004-01-2151). Warrendale, PA: SAE International. Retrieved from http://papers.sae.org/2004-01-2151/

Ingre, M., Van Leeuwen, W., Klemets, T., Ullvetter, C., Hough, S., Kecklund, G., … Åkerstedt, T. (2014). Validating and Extending the Three Process Model of Alertness in Airline Operations. *PLoS ONE*, *9*(10), e108679. http://doi.org/10.1371/journal.pone.0108679

Jensen, A., & Dahl, S. (2009). Truck drivers hours-of-service regulations and occupational health. *Work: A Journal of Prevention, Assessment and Rehabilitation*, *33*(3), 363–368. http://doi.org/10.3233/WOR-2009-0884

Jin, S., Park, S.-Y., & Lee, J.-J. (2007). Driver fatigue detection using a genetic algorithm. *Artificial Life and Robotics*, *11*(1), 87–90. http://doi.org/10.1007/s10015-006-0406-8

Kaida, K., Takahashi, M., Åkerstedt, T., Nakata, A., Otsuka, Y., Haratani, T., & Fukasawa, K. (2006). Validation of the Karolinska sleepiness scale against performance and EEG variables. *Clinical Neurophysiology*, *117*(7), 1574–1581. http://doi.org/10.1016/j.clinph.2006.03.011

Kemp, E., Kopp, S. W., & Kemp, E. (2013). Six days on the road: Will I make it home safely tonight? Examining attitudes toward commercial transportation regulation and safety. *The International Journal of Logistics Management*, *24*(2), 210–229.

Lerman, S. E., Eskin, E., Flower, D. J., George, E. C., Gerson, B., Hartenbaum, N., … Moore-Ede, M. (2012). Fatigue Risk Management in the Workplace: *Journal of Occupational and Environmental Medicine*, *54*(2), 231–258. http://doi.org/10.1097/JOM.0b013e318247a3b0

Luo, X., Hu, R., & Fan, T. (2013). The driver fatigue monitoring system based on face recognition technology. In *2013 Fourth International Conference on Intelligent Control and Information Processing (ICICIP)* (pp. 384–388). http://doi.org/10.1109/ICICIP.2013.6568102

Mallis, M. M., Mejdal, S., Nguyen, T. T., & Dinges, D. F. (2004). Summary of the Key Features of Seven Biomathematical Models of Human Fatigue and Performance. *Aviation, Space, and Environmental Medicine*, *75*(3), A4–A14.

North American Fatigue Management Program. (2012). *Module 9: Driver Scheduling and Tools*. Retrieved from http://www.nafmp.org/en/downloads.html

Rau, P. S. (2005). Drowsy driver detection and warning system for commercial vehicle drivers: field operational test design, data analyses, and progress. In *Proceedings of 19th International Conference on Enhanced Safety of Vehicles*.

Shahid, A., Wilkinson, K., Marcu, S., & Shapiro, C. M. (2011). Karolinska Sleepiness Scale (KSS). In A. Shahid, K. Wilkinson, S. Marcu, & C. M. Shapiro (Eds.), *STOP, THAT and One Hundred Other Sleep Scales* (pp. 209–210). Springer New York. Retrieved from http://link.springer.com/chapter/10.1007/978-1-4419-9893-4_47

Summary of Hours of Service Regulations. (2013, December 30). [Text]. Retrieved April 22, 2015, from http://www.fmcsa.dot.gov/regulations/hours-service/summary-hours-service-regulations

Van Dongen, H. P. A. (2004). Comparison of Mathematical Model Predictions to Experimental Data of Fatigue and Performance. *Aviation, Space, and Environmental Medicine*, *75*(3), A15–A36.

Xu, H., Chen, Z.-L., Rajagopal, S., & Arunapuram, S. (2003). Solving a Practical Pickup and Delivery Problem. *Transportation Science*, *37*(3), 347–364. http://doi.org/10.1287/trsc.37.3.347.16044

Zhang, M., Longhui, G., Wang, Z., Xu, X., Yao, B., & Zhou, L. (2014). Hybrid Model for Early
Onset Prediction of Driver Fatigue with Observable Cues. *Mathematical Problems in
Engineering*, *2014*, e385716. http://doi.org/10.1155/2014/385716

Chapter 4: An Examination of Weaknesses in Hours of Service Regulations using the Truck Driver Scheduling Problem with Fatigue Monitoring

An Examination of Weaknesses in Hours of Service Regulations using the Truck Driver Scheduling Problem with Fatigue Monitoring

Abstract

In the United States, the Federal Motor Carrier Safety Administration (FMCSA) has implemented regulations for commercial truck drivers aimed at decreasing crashes and safety incidents. Most, if not all, of these regulations focus on limiting the working and driving time of the drivers so that they will be alert while they are operating vehicles. We examine a critical shortcoming in these regulations, specifically related to assumptions made about the rest and alertness of a driver at the start of the work week.

4.0 Introduction

As we established in chapter 3, driver fatigue is a major factor in vehicle crashes. Concern for driver fatigue is elevated for commercial truck drivers who often operate trucks hauling heavy loads (up to 80,000 lbs. without a permit) at Interstate speeds. In the United States, the responsibility for regulating driving time, known as hours of service (HOS), for commercial truck drivers falls under the auspices of the Federal Motor Carrier Safety Administration (FMCSA). Technically, FMCSA HOS rules apply only to trucks that travel on Interstate highways, though practically speaking most states have the same or very similar rules which are to be enforced on other roads as well.

The FMCSA was established on January 1, 2000 "pursuant to the Motor Carrier Safety Improvement Act of 1999 (49 U.S.C. 113)" ("FMCSA," 2014). Since that time, FMCSA has released several HOS rule modifications. For the purposes of this work, we are focused only on one area of the HOS rules; specifically, the rules aimed at ensuring drivers do not begin their work week fatigued.

In 2003, FMCSA introduced a rule known as the "34 hour restart" ("Revised Hours-of-Service Rule to Help Ensure Truck Drivers Get Adequate Rest," 2003). This allowed drivers to reset their weekly HOS related hours to zero (essentially starting over), if they took 34 consecutive hours off. The goal was to help drivers get adequate rest (e.g. a weekend off) before starting a full week of driving.

In 2011, in an attempt to strengthen the effect of the 34 hour restart rule, FMCSA updated the HOS rules to require that at least two periods between 1 a.m. and 5 a.m. be included in the 34 hours off. This would theoretically increase the chances that the drivers got two good nights of

sleep prior to starting their work week. This shows that FMCSA has an interest in ensuring that drivers do not start their work week off in a fatigued state.

On July 1, 2013, FMCSA was ordered to stop enforcing the two periods between 1 a.m. and 5 a.m. ("Summary of Hours of Service Regulations," 2013). Evidently there was not enough empirical evidence to show that requiring the two periods between 1 a.m. and 5 a.m. had a positive effect on safety. FMCSA was ordered to complete a study on the effects of the change and that study (FMCSA, 2015) is currently in progress.

This paper does not aim to study the 34 hour restart rule specifically. Rather, it aims to demonstrate empirically the importance of beginning the work week in a non-fatigued state in terms of improving overall alertness of the driver throughout the week. In addition, we show that instead of making regulations that increase the likelihood that someone is rested (e.g. by requiring they not be working during specific time periods) a more effective approach would be to actually measure alertness.

We reviewed several countries and the European Union's HOS regulations and found none to have provisions related to how rested a driver needs to be when they begin their work. Solutions similar to the 34 hour restart rule are common, with varying time requirements in order to restart. This means that someone could be awake for an entire day (or longer) and immediately begin driving a fully loaded tractor trailer for another 8 straight hours before being required to take their first break yet still be in compliance with US HOS regulations. The only stipulation is that the driver not be driving, on-duty, or otherwise working for their employer during the 34 hour period prior to their work week starting. A publicized example of this occurred on June 7, 2014 in the fatal crash involving comedians Tracy Morgan and James McNair (Hanna & Marsh,

2015). We are therefore motivated to look at how different levels of fatigue at the beginning of the work week might affect the alertness of the driver for the remainder of the route.

We will leverage the Truck Driver Scheduling Problem with Fatigue Monitoring (TDSPFM) which was introduced in chapter 3. This allows us to produce schedules where the driver is likely to remain alert and thereby reduce the likelihood of being involved in a crash. The TDSPFM also incorporated HOS constraints and allows us to measure alertness at all stages along the route.

In this research, we examine the effects of different rest and alertness levels at the beginning of the work week. We show how this should be considered a key weakness in existing HOS regulations and how these starting levels can limit the effectiveness of alertness thresholds in the TDSPFM. Additionally, we show how using the TDSPFM at the planning level can help identify times in the schedule where the driver may be vulnerable to fatigue and how modifications to the TDSPFM can allow for verifiable alertness scores.

4.1 Literature Review

Fatigue detection among drivers is an active area of research and product development (Brown, Johnson, & Milavetz, 2013; Chang & Chen, 2014; Luo, Hu, & Fan, 2013; Zhang et al., 2014). In all of the articles we reviewed, the focus was on detecting driver fatigue in the vehicle. Generally speaking, the goal in existing research is to alert the driver before an accident occurs by using data captured during the trip. For instance, one could use yawn and blinking frequency to predict fatigue using computer vision technologies (Jin, Park, & Lee, 2007) and then alert the driver when the fatigue estimate crosses some pre-defined threshold.

We contend that near real-time fatigue detection is a valuable safety net but that additional planning measures should try to reduce the need for such detection as much as possible for both cost and safety reasons. From a safety standpoint, the real-time detection of fatigue requires the presence of either fatigue or some characteristics attributable to the imminent onset of fatigue. Eliminating the fatigue in the first place would produce more alert and, by definition, less fatigued drivers. Considering costs, if we assume a driver takes an unplanned break when fatigue is detected, that will impact not just the current leg of the trip but potentially all stops for the remainder of the route. Factoring in time windows, which are almost always a real-world constraint, this becomes a much more important consideration to address. Should any stops along the route have tight deadlines, unplanned stops (particularly unplanned rest stops which are often 8 hours or longer) may result in the adjusted schedule violating (at least) the time window constraint. The earlier in the schedule the unplanned stop happens, the greater the likelihood of the change having compounding effects on other stops in the route.

4.1.1 Three Process Model of Alertness

The Three Process Model of Alertness (TPMA) model utilized in this research consists of three primary processes that have been previously published (Åkerstedt, Folkard, & Portin, 2004) and are described briefly below. Process C represents the circadian influence on alertness; this process encompasses the effect that the time of day can have on sleepiness. Process S describes the exponential decline in alertness as a function of the time awake. This decline is then reversed in process S' which describes the recovery as a function of the time asleep. Figure 4-1 below shows the way these processes effect alertness at different times during the day. Process W describes the lack of alertness at the time of waking up. Since we will not assume driving takes place immediately following waking up, we ignore the W component of the TPMA

model for this research. Finally, since the original publication of the TPMA model, another process has been added (Åkerstedt, Axelsson, & Kecklund, 2007). This process is U, which stands for "ultradian" and explains an afternoon dip in alertness.
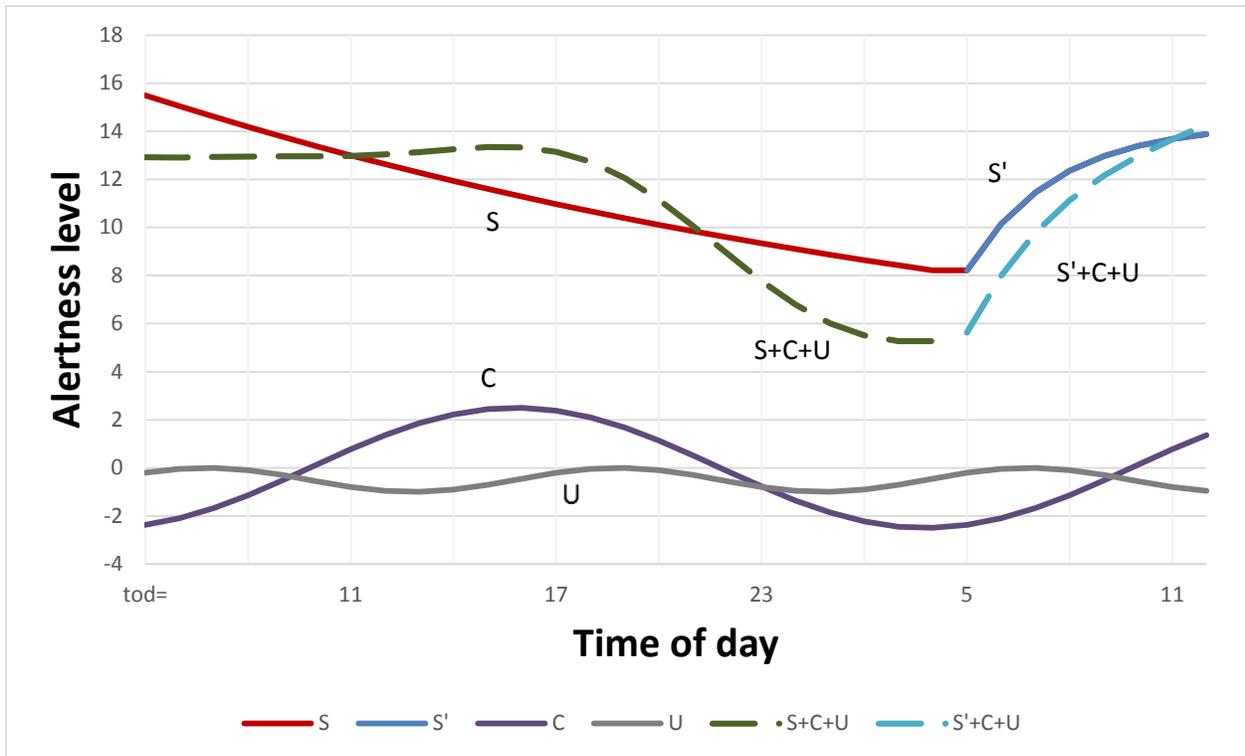

Figure 4-1: Components of the TPMA

For the implementation of our TDSPFM model, we will use the validated TPMA model and parameters presented in (Ingre et al., 2014). Therefore, our model predicts alertness while driving as the summation of S+C+U. This produces an alertness score with values ranging from 1 to 21. According to the sleep literature, a TPMA value (or score) of "3" corresponds to extreme sleepiness, while "14" represents high alertness, and "7" to a borderline sleepiness threshold

(Åkerstedt et al., 2004). While a minimum allowed alertness score is a required parameter in our model, the appropriate value of this parameter is a matter for further research.

4.1.2 Alertness Score Levels

An advantage of the TPMA alertness score is that it can be transformed into a prediction of the subjective Karolinska Sleepiness Scale (KSS) which has been shown to be a valid means for measuring sleepiness (Kaida et al., 2006). The KSS is a scale that "measures the subjective level of sleepiness at a particular time during the day" (Shahid, Wilkinson, Marcu, & Shapiro, 2011) and was introduced by Åkerstedt and Gillberg (1990). Since the KSS can be obtained simply by asking a driver to rate their sleepiness on a 9 point Likert scale, future research could validate the schedules and alertness scores proposed by our model.

The transformation from TPMA to KSS is based on the work by Ingre et al. (2014) and can be calculated as follows:

KSS = 9.68 – (0.46 * TPMA alertness score) + 1.07

This allows us to readily transform a TPMA alertness score into a KSS score and vice-versa. More research has been done correlating KSS to driver drowsiness and suggests that an appropriate sleepiness threshold for drivers may be higher than the suggested value of 7 from existing sleep literature (Åkerstedt, Connor, Gray, & Kecklund, 2008; Ingre et al., 2006; Liu, Hosking, & Lenné, 2009).

One important difference between the TPMA alertness scale and the KSS is that the TPMA measures alertness whereas the KSS measures sleepiness. Therefore, a high KSS value implies that that subject is sleepy and thus less alert whereas a high TPMA alertness score

implies the subject is highly alert. To complicate matters, both the KSS and the TPMA use a value of 7 as the sleepiness threshold. This can make switching between TPMA and KSS confusing. To avoid confusion, we only use TPMA values for results and parameter values, converting KSS values from previously published research to TPMA equivalents.

4.1.3 Sleep Assumptions

As we did in chapter 3, we will remain conservative with some of the assumptions we make regarding the quality and length of sleep a driver will get in a week. We make the assumption that when long rest breaks are taken, the driver does in fact use that time to get good, recovering sleep. We do not factor in the stochastic nature of sleep, whereby many external factors (such as noise, light levels, cabin temperatures, etc.) can ultimately affect the total amount of sleep one is able to attain. Finally, we do not consider caffeine or other drug use that could affect driver alertness.

4.1.4 Hours of Service Regulations

In general, HOS regulations function primarily by imposing rules related to how long drivers can stay on the road, the conditions relating to rest frequency, and the duration/types of rest periods. As a result, a drivers schedule for a particular day might look like Figure 4-2 below.

Figure 4-2: A Truck Driver Schedule (modified from Goel, 2012)

For the purposes of the remainder of this paper, we will consider the HOS regulations in the United States. Specifically we will use the parameters defined and explained in Goel (2014). While we use the US regulations, other HOS restrictions could easily be accommodated by our model.

4.2 TDSPFM Model

An extensive presentation and description of our model implementation can be found in chapter 3. Below, we will cover in brief the relevant variables from the previously presented model and focus on how our implementation for this chapter differs.

We consider a sequence of N locations which will be visited by a truck driver. Each location $i \in N$ has a time window and some duration of work associated with it. We let the opening time window of each location be zero for the sake of simplicity in producing results and just focus on the closing time window of L_i for each $i \in N$. The work duration is denoted as w_i for each $i \in N$. We let $t_{i,i+1}$ be the amount of time it takes to drive between locations i and the next location. The arrival time and departure time of each location i is denoted as A_i and D_i respectively.

Because the sequence of locations in our model is fixed, our focus is on choosing the optimal schedule of rest times at each location in order to minimize duration. Rest times are denoted below as r_i for each $i \in N$. The duration we wish to minimize can be calculated as the difference between the arrival time at the final location and the departure time of the first location.

We chose to use parameters representative of the HOS regulations in the United States. These parameters and their values used for our implementation of the TPMA are displayed in

Table 4-1 below. The parameter values for the TPMA implementation come from Ingre et al.

(2014).

Parameter Name	Value (h)	Description
t^{rest}	10	Minimum rest time to be considered a long break
t^{break}	.5	Minimum Break
$t^{horizon}$	168	Planning Horizon (1 week)
t^{drive}	11	Maximum drive time since last break
$t^{elapsed|R}$	14	Maximum time since last long rest
$t^{elapsed|B}$	8	Maximum time since last break
$t^{maxsleep}$	10	Maximum amount of sleep allowed
$t^{awakedelay}$	0.5	The delay after waking before driving can begin
$t^{sleepdelay}$	1	The delay after rest begins before falling asleep
t^{max_rest}	16	Maximum rest period allowed
la	2.4	Lower asymptote of the internal alertness scale
d	-0.0353	Decay in alertness
g	-0.38135645	Recovery multiplier
ha	14.3	Higher asymptote of the internal alertness scale
bl	12.5	Break level of recovery function S'
Cm	0	Mesor of process C
Ca	2.5	Amplitude of process C
p	16.8	Default circadian phase
Um	0.5	Mesor of process U

Ua	0.5	Amplitude of process U

Table 4-1: TDSPFM Parameters

In order to abide by the HOS regulations, we need to keep track of the following variables which are presented below along with their associated equations:

- Time driving upon arriving at location i, dt_i:

$$dt_i = \begin{cases} dt_{i-1} + t_{i-1,i}, & r_{i-1} < t^{rest} \\ t_{i-1,i} & , & r_{i-1} \geq t^{rest} \end{cases}, \forall i \in N$$

- Time working since last long rest upon arriving at location i, st_i:

$$st_i = \begin{cases} st_{i-1} + t_{i-1,i} + w_i, & r_{i-1} < t^{rest} \\ t_{i-1,i} + w_i & , & r_{i-1} \geq t^{rest} \end{cases}, \forall i \in N$$

- Time since last break upon arriving at location i, stb_i:

$$stb_i = \begin{cases} stb_{i-1} + t_{i-1,i} + w_i, & r_{i-1} < t^{break} \\ t_{i-1,i} + w_i & , & r_{i-1} \geq t^{break} \end{cases}, \forall i \in N$$

To determine the alertness of the driver at a given location i ($alertness_i$), we need to introduce the following variables:

- The time awake upon arriving at location i, taw_i:

$$taw_i = \begin{cases} taw_{i-1} + t_{i-1,i} + r_{i-1} + w_i & , & r_{i-1} < t^{rest} \\ t_{i-1,i} + (r_{i-1} - t^{maxsleep}) + t^{awakedelay} + w_i, & r_{i-1} \geq t^{rest} \end{cases}, \forall i \in N$$

- The time of day upon arriving at location i, tod_i:

$$tod_i = A_i \bmod 24 \ \forall i \in N$$

Finally, we show the calculations related to the components of the TPMA which will ultimately allow us to compute an alertness score at a given location i (*alertness_i*):

- Process S at location i, S_i:

$$S_i = la + (S'_{i-1} - la) * e^{d*taw_i} \ \forall \ i \in N$$

- Subprocess SB_i, used to determine the proper S'_i to use:

$$SB_i = ha - (ha - ss)e^{g(t^{maxsleep} - t^{sleepdelay})} \ \forall \ i \in N$$

- Process S' at location i, S'_i:

$$S'_i = \begin{cases} ha - (ha - bl)e^{g(t^{maxsleep} - t^{sleepdelay} - bt_i)} & , & r_{i-1} \geq t^{rest} \ AND \ SB_i \ \geq bl \\ S_i + g(t^{maxsleep} - t^{sleepdelay}) * (bl - ha), & r_{i-1} \geq t^{rest} \ AND \ SB_i \ < bl \\ S'_i & , & r_{i-1} < t^{rest} \end{cases} , \forall i \in N$$

- Process C at location i, C_i:

$$C_i = Cm + Ca * \cos\left(\left(2\frac{\pi}{24}\right) * (tod_i - p)\right) \ \forall \ i \in N$$

- Process U at location i, U_i:

$$U_i = Um + Ua * \cos\left(\left(2\frac{\pi}{12}\right) * (tod_i - p - 3)\right) \ \ \forall \ i \in N$$

The above processes allow us to compute the alertness score at a given location i (*alertness_i*):

$$alertness_i = S_i + C_i + U_i, \quad \forall \ i \in N$$

Our formulation of the TDSPFM is then given as follows,

Minimize:

$$A_{last} - D_{first} \tag{1}$$

Subject To:

$$A_i + r_i + w_i = D_i \quad \forall i \in N \tag{2}$$

$$D_i + t_{i,i+1} = A_{i+1} \quad \forall i \in N \tag{3}$$

$$dt_i \leq t^{drive}, \forall i \in N \tag{4}$$

$$st_i \leq t^{elapsed|R}, \forall i \in N \tag{5}$$

$$stb_i \leq t^{elapsed|B}, \forall i \in N \tag{6}$$

$$A_{last} + W_{last} \leq t^{horizon} \tag{7}$$

$$0 \leq r_i \leq t^{max\_rest}, \quad \forall i \in N \tag{8}$$

$$A_i \leq L_i, \quad \forall i \in N \tag{9}$$

$$minalertness_i \geq TPMA^{min}, \quad \forall i \in N \tag{10}$$

$$r_i \geq 0, \quad \forall i \in N \tag{11}$$

The objective in (1) is to minimize the duration of our route from our first location until

we arrive at the last location. Constraints (2) and (3) ensure our arrival and departure times

capture the time consumed at each location. Constraints (4)-(7) enforce the HOS regulations. Constraint (4) is the HOS drive time constraint. Constraint (5) shows the HOS working time constraint. Constraint (6) ensures that a short rest break is taken in accordance with HOS regulations. Constraint (7) ensures that the arrival time and associated work time of the last location (A_{last}, w_{last}) occur within the time horizon specified for the problem. We also set a maximum rest time ($t^{\max\_rest}$) and constraint (8) ensures that the rest time at each location *(r$_i$)* is within that limit.

Constraint (9) is the time window restriction ensuring that the arrival time of each location is prior to the time window of that location (*L$_i$*) closing. For the sake of simplicity, we allowed an opening time window at each location to be 0 though an additional constraint to account for an opening time window can easily be added.

Constraint (10) allows us to specify an alertness threshold ($TPMA^{\min}$) and ensure that the alertness score stays above the specified alertness threshold. Should we only want to use the hours of service constraints, we set $TPMA^{\min} = 0$. Finally, constraint (11) ensures the rest time at location *i* (*r$_i$*) is non-negative.

The TDSPFM is a planning problem, thus we are more concerned with finding good solutions than fast run-times. Therefore, we coded this model in Microsoft Excel and solved it using the Evolutionary solution method of the built-in Solver.

Since we already know the travel times and sequence of locations to be visited, we coded a repair function that would force a rest of at least t^{rest} in length if continuing would violate one of the hours of service restrictions. As is common in GA formulations, we used penalty functions to capture violations of time windows, HOS regulations, and alertness constraints.

4.3 Example Problem

Using our TDSPFM model in combination with an actual truck driver's schedule taken from his "Driver's Daily Log Book", we can observe the values of key performance indicators pertaining to fatigue. We can also observe how changing starting alertness levels affect these key values.

After contacting a truck driver, we received two of his log books and randomly chose one week to model. The driver's typical work day was from 1 a.m. until noon. For the weekly route we modeled, the duration of the route was 125 hours and involved 30 stops (including the beginning and ending locations) with varying durations per stop.

We assume that the driver was awake for 2 hours at the start of his work shift and prior to driving and that the initial TPMA alertness level was set at 12.5, which corresponds to a KSS score of 5 (neither sleepy nor alert). In this case, the minimum alertness for the driver's schedule was 7.11 with an average alertness of 8.48. Recall from above, the generally accepted alertness threshold in the TPMA literature is 7, so the minimum alertness score is slightly above the published alertness thresholds and the average score is uncomfortably close to that threshold.

Next, we adjusted the starting alertness level to investigate what effect that had on the key performance indicators. The results are displayed in Table 4-2 below.

Initial Alertness	Minimum Alertness	Average Alertness
High - 12.5	7.11	8.48
Medium - 10.32	7.01	7.96
Low - 8.15	5.89	7.37

Table 4-2: Example Schedule Alertness Levels

Graphically, the alertness levels are represented in Figure 4-3. As the graph shows, the low initial starting alertness reduces the average alertness over the example weekly schedule. Additionally, the minimum alertness is substantially reduced when the initial alertness level is low.
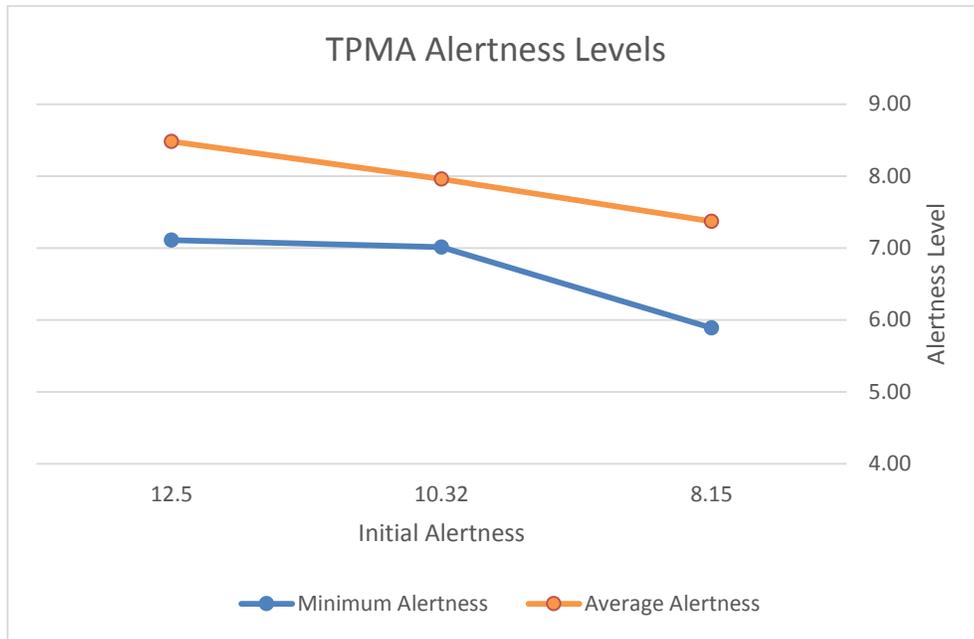

Figure 4-3: TPMA Alertness Levels

To gain an understanding of the alertness level over the course of the week, we present Figure 4-4 below. We graph the different components of the TPMA as a function of the different stops in the example week's route using a starting alertness of 12.5. Rests are taken at points where the S curve slopes upward. However, the resulting recovery can be dampened by other components in the TPMA, particularly the circadian effect (process C). Once can also observe how on several occasions the minimum alertness level gets close to the TPMA threshold of 7.
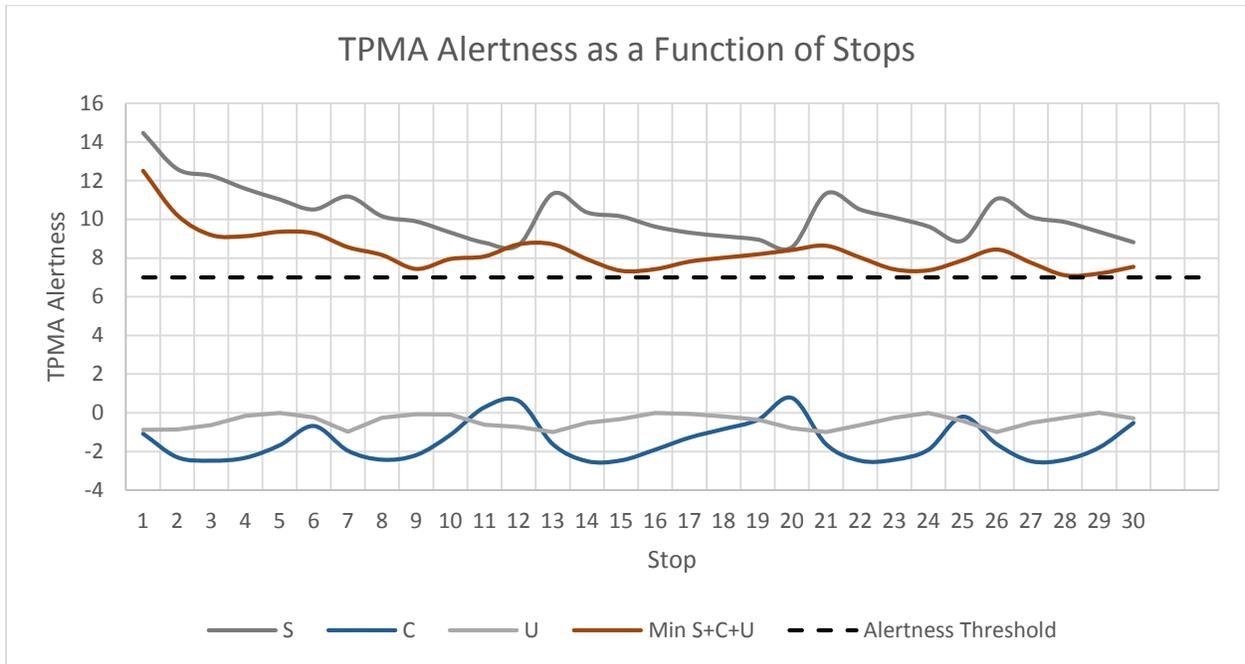
Figure 4-4: TPMA Alertness in Example Schedule

The alertness measures captured from the sample log entries show that the less alert the driver is when starting the week, the less alert the driver will be (both in the overall minimum and on average) throughout the week. This result shows that initial fatigue impacts the entire week and is not something a driver recovers from after the first rest break. However, since this is only the results from analyzing one driver's schedule, we cannot conclude that this would be true for other drivers and schedules and thus need further computational testing.

4.4 Computational Testing and Results

To get a better understanding of the magnitude of the effect of the starting alertness level, we created 35 benchmark problems in the same manner as described in chapter 3 and allow for varying the starting alertness level. We used the following starting alertness levels:

- High, alertness level 12.5

- Medium, alertness level 10.32

- Low, alertness level 8.15

We then use the same minimum alertness thresholds ($TPMA^{min}$) as we used in chapter 3 above:

- not tired, $TPMA^{min} = 9.24$

- semi-tired, $TPMA^{min} = 8.15$

- tired, $TPMA^{min} = 7.07$

- hours of service only, $TPMA^{min} = 0$

We study the combinations of the above configurations, which leads to 10 sets of 30 problems from which to derive our results. (Note there would be no feasible solutions available with starting TMPA levels of 8.15 and required $TPMA^{min}$ of either 7.07 or 8.15.) The box and whiskers plots in Figures 4-5, 4-6, and 4-7 below show the results of these sample problems. We group the problems along the x-axis by starting alertness level.

Figure 4-5 shows how each problem set performed with respect to the duration of the routes. These results are consistent with the results presented in chapter 3. At the "not tired" level ($TPMA^{min} = 9.24$) we see a slight increase in duration and also an increase in the variability of the result. Overall, the starting alertness level has little effect on the duration of the schedule.
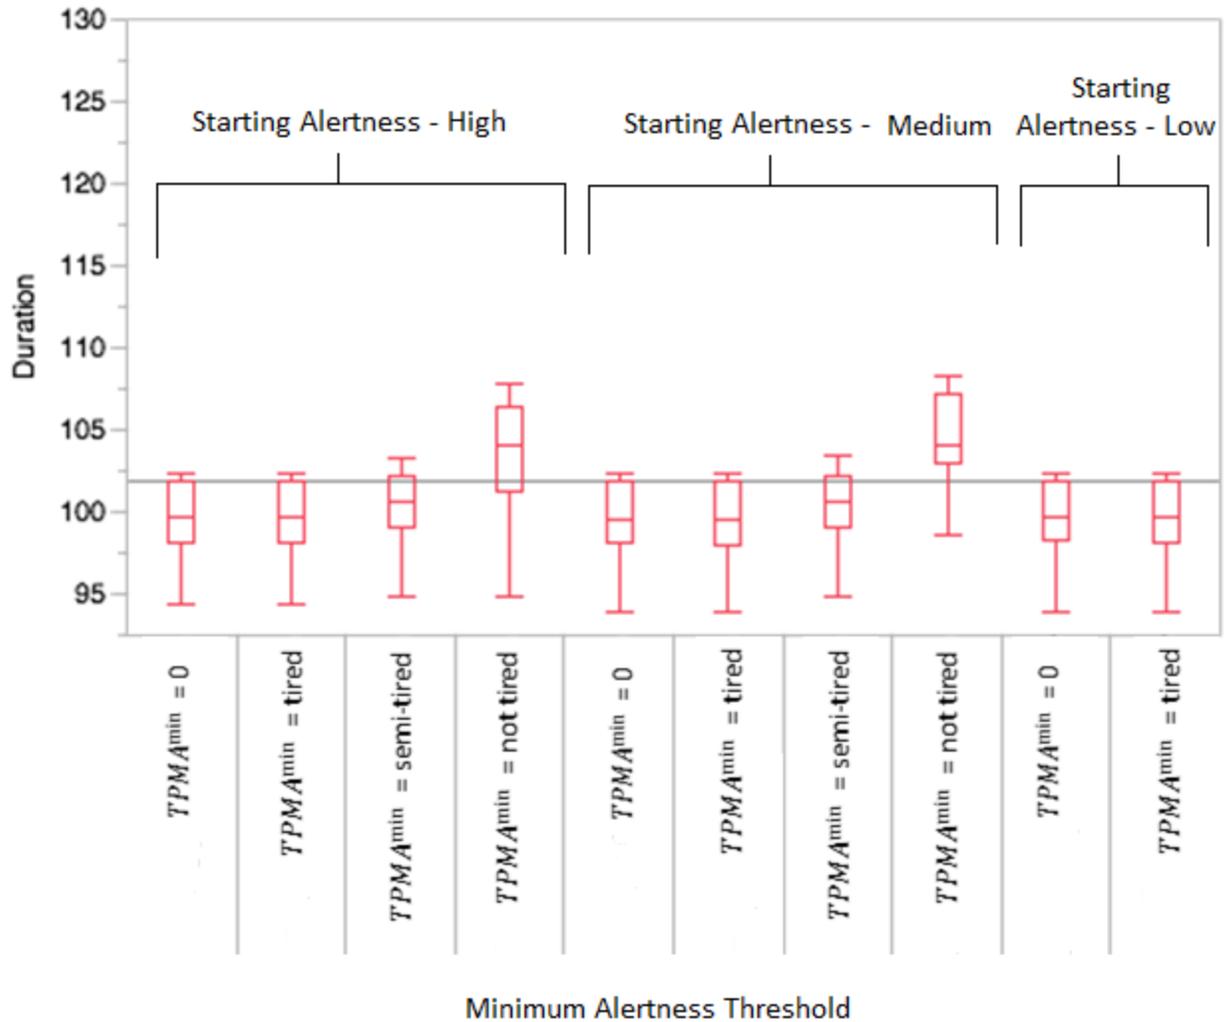
Figure 4-5: Box and Whiskers Plot - Duration

In Figure 4-6, we present the results of the minimum alertness for our problem sets. The minimum alertness observations themselves are in line with the setting of $TPMA^{min}$ in that the lower whisker of the plot is roughly at the $TPMA^{min}$ value. The most interesting result from this figure is the observation that as $TPMA^{min}$ increases, the variability of the result decreases. In general though, the results show the same behavior presented in Table 3-2 of chapter 3. As far as the minimum alertness of the trip, the initial starting alertness level has little effect if we throw out the obvious infeasible solutio
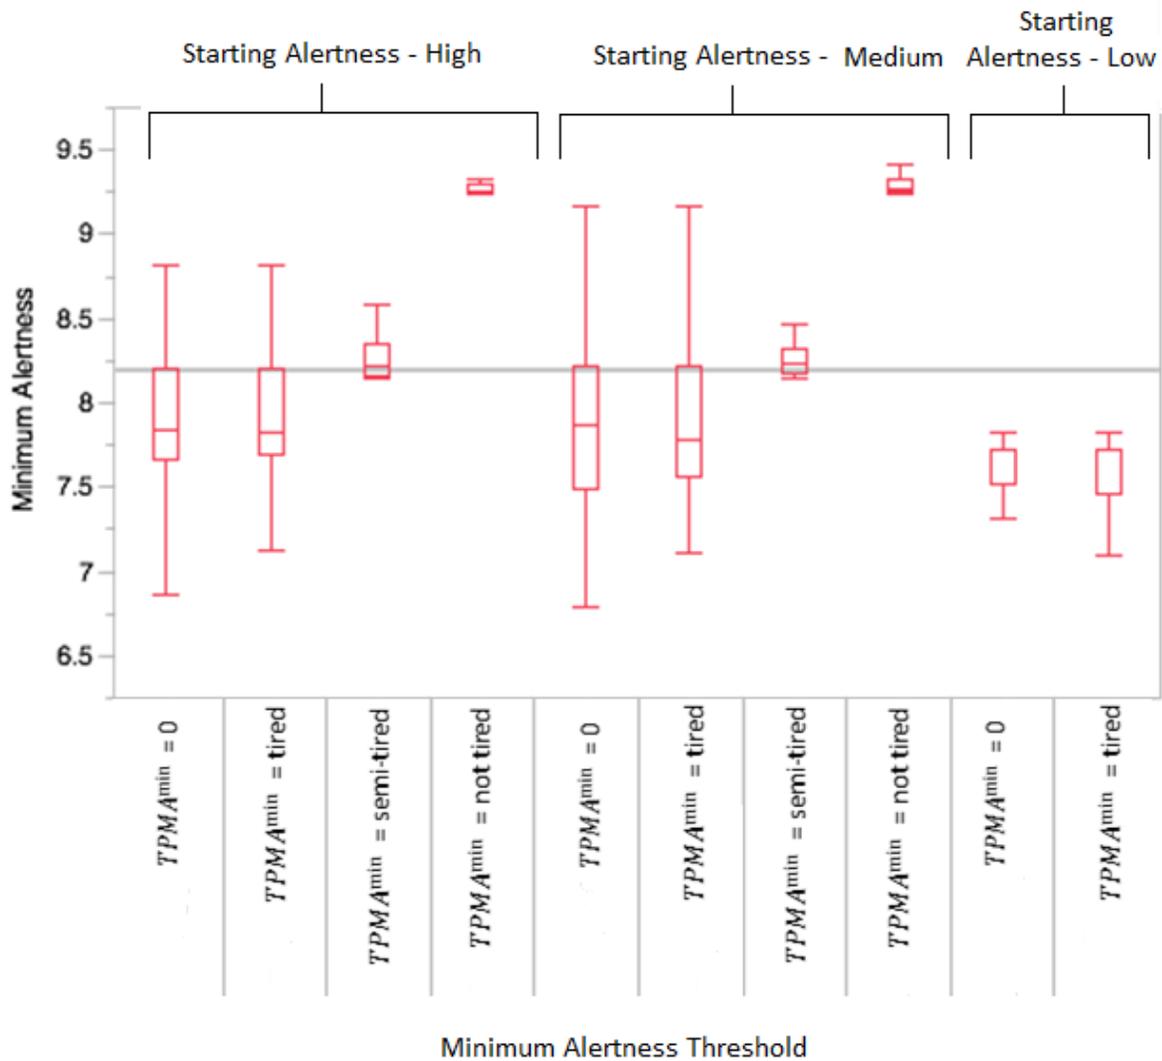
Figure 4-6: Box and Whiskers Plot – Minimum Alertness Level

The most interesting observation derived from our results is presented in Figure 4-7, where we compare the averaged alertness level of our problem sets. Here we can clearly see that the starting alertness level has an impact on the week's average alertness level. The impact is most pronounced when the starting alertness is low.
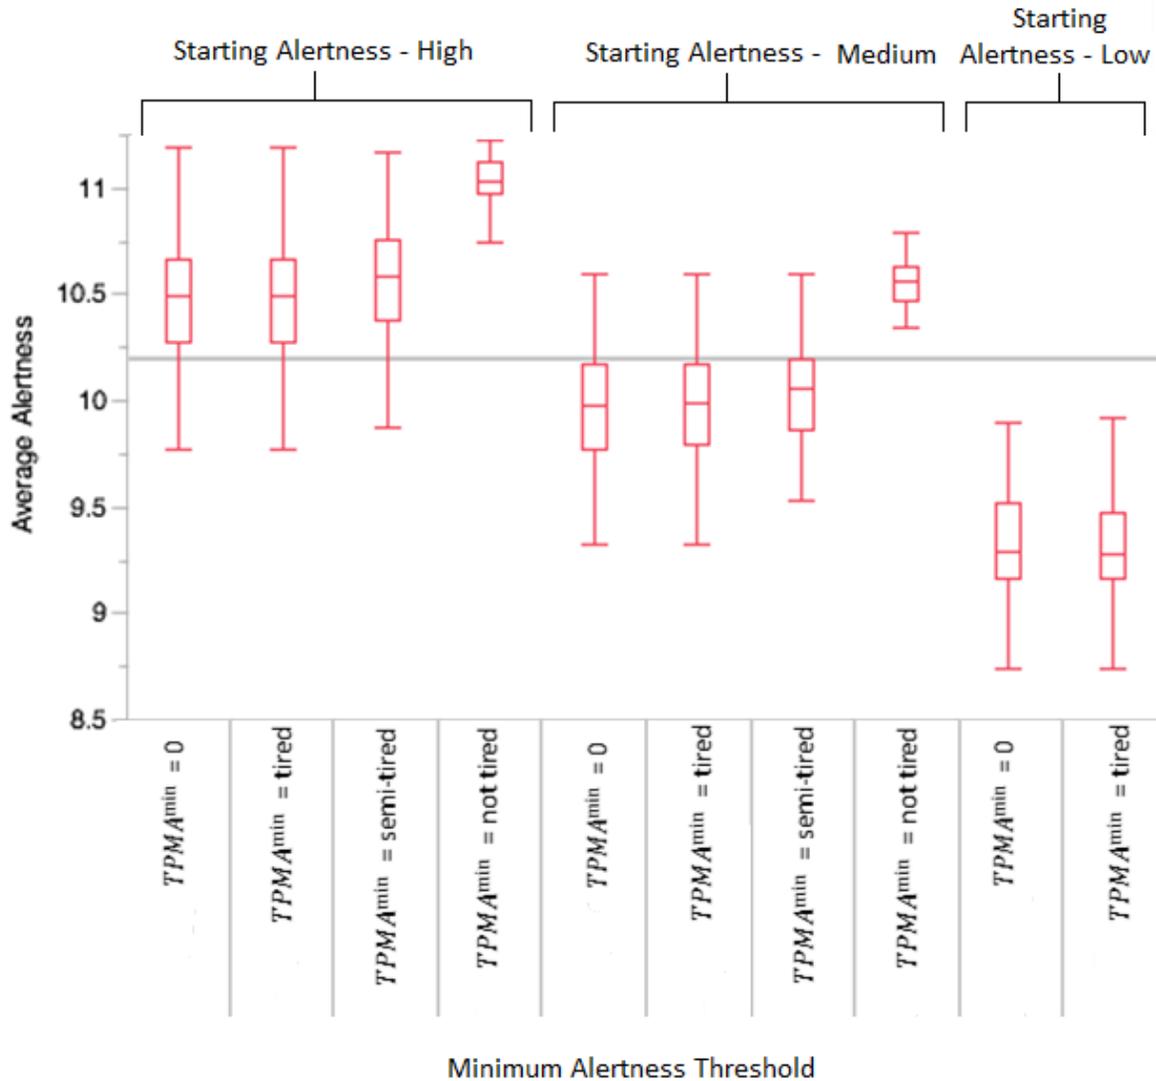
Figure 4-7: Box and Whiskers Plot – Average Alertness Level

Much of the focus on the research stream presented thus far has been on the minimum alertness level. In other words, trying to keep the driver from becoming too fatigued at any given point along the route in order to avoid a crash or other disruptive event at that given point in time. This is a valid idea to study and we feel it has merit as a possible crash reduction or fatigue management strategy.

However, while the minimum alertness covers a single point in time along the route, the average alertness looks more inclusively at the time driving during the week as a whole. Presumably there will be many more situations where the driver is at or near the average alertness level where they will need to make important decisions that could be impaired by fatigue. Therefore, as a fatigue management strategy, the average alertness may be equally or perhaps more important than the minimum alertness level.

Average alertness becomes even more important when we want to look at driver decision making as a whole as opposed to considering situations where the driver falls asleep at the wheel or begins to lose control of the vehicle due to fatigue. The later may very well occur at the minimum alertness level of the trip. When considering maneuvers that involve the driver's decision making, such as making a maneuver to avoid a potential crash or the need to make quick reactions, the average alertness level is a more appropriate measure. It is plausible to conjecture that the more alert the driver is on average, the more likely it is that they will have the capacity to make good decisions and to make them quickly.

4.4.1 Highway To The Danger Zone

To investigate this further, we seek to set an alertness level where external factors (poor sleep, schedule change, etc.) would have an increased chance at bringing the driver's alertness below the $TPMA^{\min}$. We will refer to this level as the "danger zone" and we are particularly interested in the periods of time driving when the driver's alertness level is in this danger zone. To allow us to compare results across model configurations, we will use a danger zone value of 9.7 (a 5% increase in the "not tired" $TPMA^{\min}$ parameter value). However, we acknowledge that the appropriate value of this threshold is a topic of research which we hope to investigate further in the future.

This danger zone would be the areas in the schedule where the driver's TPMA alertness level is between 9.7 and the $TPMA^{\min}$ parameter value. From a fatigue management perspective, this becomes an important subset of the schedule to be concerned with motivated primarily by real-world scheduling uncertainties and assumptions. For instance, the TDSPFM assumes that a driver spends most of their long rest break getting good, recoverable sleep. On points along the route where the driver's predicted alertness level will be close to the $TPMA^{\min}$, it is critical that the prior sleep periods did result in the modeled alertness recovery. Should something happen that results in the driver getting less sleep than the model predicted, the alertness recovery would be reduced and the resulting alertness on later segments of the route may end up being below the $TPMA^{\min}$.

To see a graphical representation of our proposed danger zone, consider Figure 4-8 below. We use one of the example problems created above to model the minimum TPMA alertness score at each stop on the optimal schedule. Figure 4-8 allows us to compare alertness scores between different model parameters (no $TPMA^{\min}$ with a starting alertness of Low compared to $TPMA^{\min}$ = "not tired" and a starting alertness of High) and observe the portion of the schedule spent in the danger zone (below the red line).
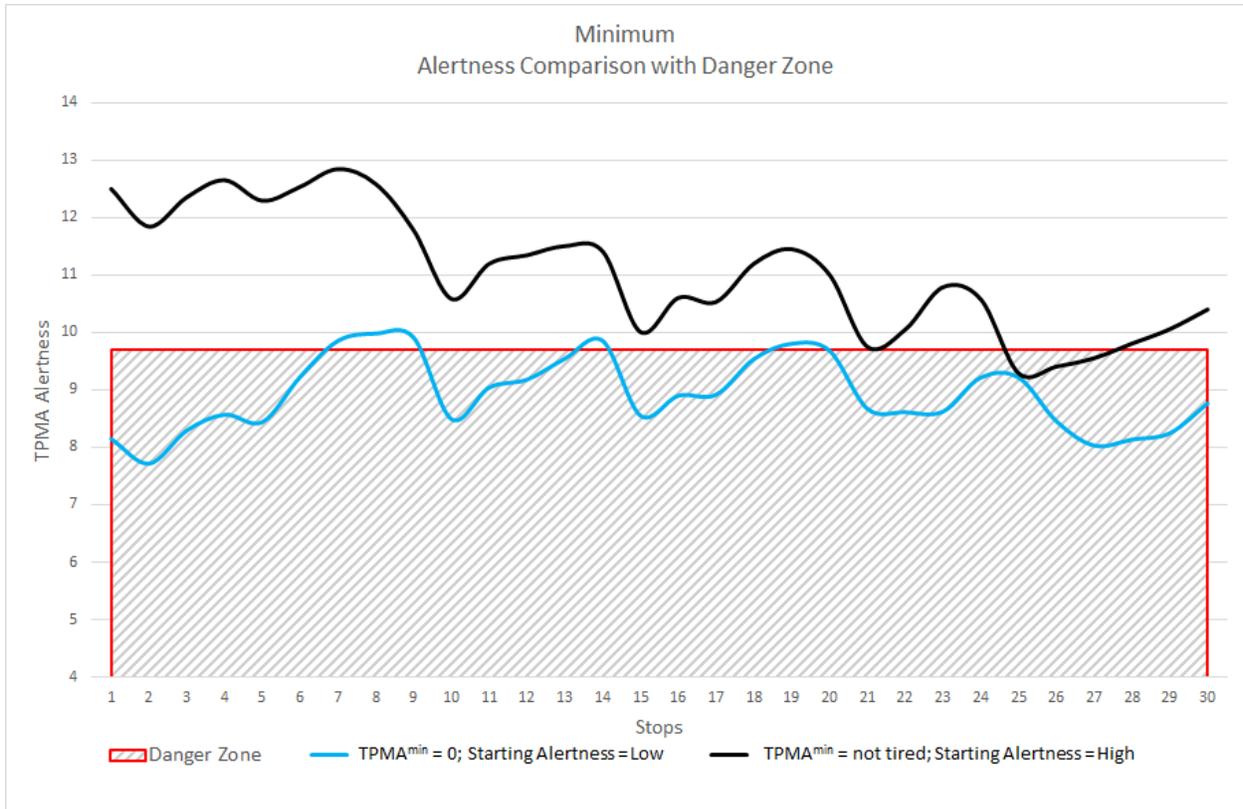
Figure 4-8: Minimum Alertness Comparison

For a more rigorous analysis, Figure 4-9 below shows the modeled percentage of time driving while the alertness level is in the danger zone across all of our example problems.
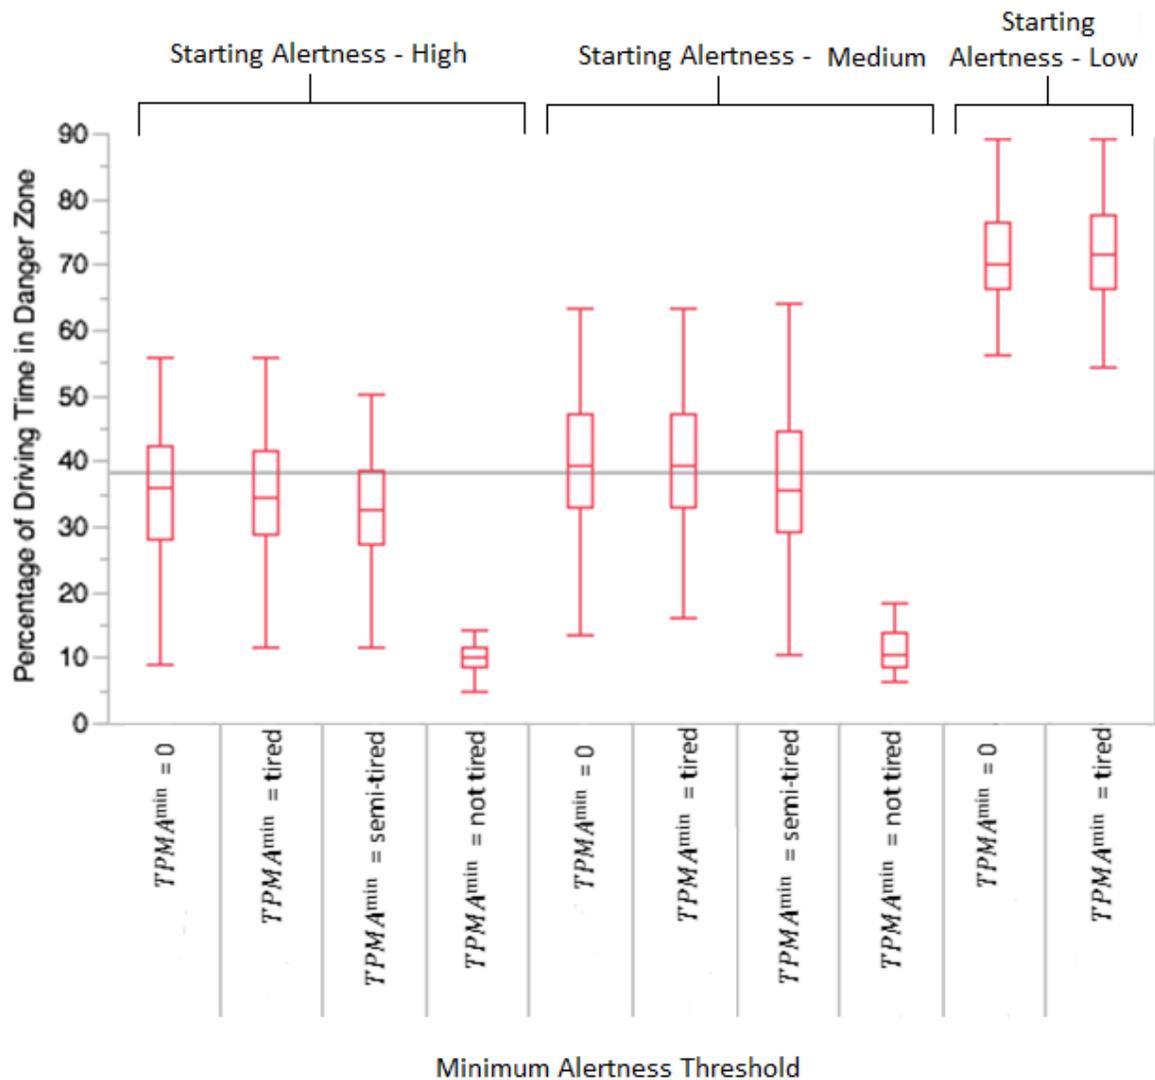
Figure 4-9: Box and Whiskers Plot – Percentage of Time Driving in Danger Zone

Intuitively, when $TPMA^{min}$ is set to "not tired" (9.24), the portion of the danger zone that is in the feasible solution space is smaller and thus percentage of time is lower. However, the more important observation is how much of an effect starting alertness has on the percentage of time in the danger zone. Starting at the lower level of alertness results in a median of 70% of time driving in the danger zone. This is also supported in Figure 4-7 which shows that when starting at the lower level of alertness the median average alertness is in the danger zone.

In summary, we conclude that the average alertness and percentage of time in the danger zone are significantly affected by the initial starting alertness level. At the lowest level of starting alertness that we looked at, this was especially pronounced. Given that there currently exist no regulatory provisions pertaining to the initial alertness level, our research points (at least) to the need for investigation into the practicality of such a provision.

Even though existing regulations may not address initial alertness specifically, our work shows its effect on the overall alertness levels of a driver throughout their work week. Additionally, our work shows how to calculate and predict alertness levels. This serves as a solid starting point for companies to take proactive measures to prevent drivers from beginning work in a compromised alertness condition and to monitor their alertness throughout the week. In light of our work and results presented above, companies could face significant liability issues should they neglect to take such proactive measures even in absence of regulations.

4.5 Conclusions and Future Work

Based on these results, we predict that the "Commercial Motor Vehicle Driver Restart Study" (FMCSA, 2015) will show that two overnight rests in the restart period will result in more alert drivers. This assumes that the data also shows that two overnight rests in the restart period do in fact lead to higher alertness levels at the start of the work week.

We can validate our model and prediction by using data sets obtained as part of Federal Motor Carrier Safety Administration studies such as the "Commercial Motor Vehicle Driver Restart Study" (FMCSA, 2015). Additionally, we could incorporate a study of driver schedule plans as part of a naturalistic driving study at the Virginia Tech Transportation Institute (VTTI).

Additionally, we can further investigate the appropriate value of where the upper bound on the danger zone should be. It is possible that this value could be determined as more datasets are made available for analysis. It is also quite possible that other circumstances can influence the appropriate value of the danger zone whether this be characteristics about individuals, types of roads travelled, or even different settings based on the time of day.

There are three primary areas we have identified for potential expansion of the TDSPFM model: sleep variability, sleep prediction, and driver specific alertness modifications. We can look at sleep variability from both the standpoint of individual variations and specific rest period variations. In our current implementations presented below, we look at the typical/average sleep functions. The TPMA however can easily support user specific functions; for instance, certain people may recover during sleep at different rates or be more affected by being awake and needing to work at 3 am.

Sleep prediction is an existing feature that can be implemented in the TPMA (Åkerstedt et al., 2008). This allows one to predict sleep schedules and to factor those results into the alertness recovery process (S'). We chose not to implement this feature in the research presented here; however, it can be done and it can also support individual variations as described above.

Finally, our models assume non-stochastic sleep times during rest periods. However, the amount of time that you actually sleep on a given night is dependent upon many factors (noise levels, comfort, temperature, etc.). We could incorporate a stochastic element into the model to represent how these unknown factors affect sleep duration.

As a planning problem, we should look for areas where a decision support system (DSS) could aid in selecting safer routes for drivers. This could allow for selection of parameter values

such as a minimum acceptable alertness level, or it could also alert schedulers to specific points along the route where alertness is predicted to be in the danger zone. This could allow some specific action to be taken; for example, check in with the driver around that time or take extra precautions to increase the chances that the driver got a good night's rest the night prior. Other extensions exist such as allowing for different alertness levels depending on the cargo being transported or the areas through which the route passes.

References

Åkerstedt, T., Axelsson, J., & Kecklund, G. (2007). Individual validation of model predictions of sleepiness and sleep hours. *Somnologie - Schlafforschung Und Schlafmedizin*, *11*(3), 169–174. http://doi.org/10.1007/s11818-007-0315-7

Åkerstedt, T., Connor, J., Gray, A., & Kecklund, G. (2008). Predicting road crashes from a mathematical model of alertness regulation—The Sleep/Wake Predictor. *Accident Analysis & Prevention*, *40*(4), 1480–1485. http://doi.org/10.1016/j.aap.2008.03.016

Åkerstedt, T., Folkard, S., & Portin, C. (2004). Predictions from the Three-Process Model of Alertness. *Aviation, Space, and Environmental Medicine*, *75*(3), A75–A83.

Åkerstedt, T., & Gillberg, M. (1990). Subjective and objective sleepiness in the active individual. *The International Journal of Neuroscience*, *52*(1-2), 29–37.

Brown, T., Johnson, R., & Milavetz, G. (2013). Identifying Periods of Drowsy Driving Using EEG. *Annals of Advances in Automotive Medicine*, *57*, 99–108.

Chang, T.-H., & Chen, Y.-R. (2014). Driver fatigue surveillance via eye detection. In *2014 IEEE 17th International Conference on Intelligent Transportation Systems (ITSC)* (pp. 366–371). http://doi.org/10.1109/ITSC.2014.6957718

FMCSA. (2014, March 31). [Text]. Retrieved September 21, 2015, from https://www.fmcsa.dot.gov/mission/about-us

FMCSA. (2015, April 28). Commercial Motor Vehicle Driver Restart Study [Text]. Retrieved September 1, 2015, from http://www.fmcsa.dot.gov/safety/research-and-analysis/commercial-motor-vehicle-driver-restart-study

Goel, A. (2012). The minimum duration truck driver scheduling problem. *EURO Journal on Transportation and Logistics*, *1*(4), 285–306. http://doi.org/10.1007/s13676-012-0014-9

Goel, A. (2014). Hours of service regulations in the United States and the 2013 rule change. *Transport Policy*, *33*, 48–55. http://doi.org/10.1016/j.tranpol.2014.02.005

Hanna, J., & Marsh, R. (2015, August 11). NTSB: Trucker's fatigue likely led to Tracy Morgan crash - CNN.com. *CNN*. Retrieved from http://www.cnn.com/2015/08/11/us/ntsb-tracy-morgan-crash/index.html

Ingre, M., Åkerstedt, T., Peters, B., Anund, A., Kecklund, G., & Pickles, A. (2006). Subjective sleepiness and accident risk avoiding the ecological fallacy. *Journal of Sleep Research*, *15*(2), 142–148. http://doi.org/10.1111/j.1365-2869.2006.00517.x

Ingre, M., Van Leeuwen, W., Klemets, T., Ullvetter, C., Hough, S., Kecklund, G., … Åkerstedt, T. (2014). Validating and Extending the Three Process Model of Alertness in Airline Operations. *PLoS ONE*, *9*(10), e108679. http://doi.org/10.1371/journal.pone.0108679

Jin, S., Park, S.-Y., & Lee, J.-J. (2007). Driver fatigue detection using a genetic algorithm. *Artificial Life and Robotics*, *11*(1), 87–90. http://doi.org/10.1007/s10015-006-0406-8

Kaida, K., Takahashi, M., Åkerstedt, T., Nakata, A., Otsuka, Y., Haratani, T., & Fukasawa, K. (2006). Validation of the Karolinska sleepiness scale against performance and EEG variables. *Clinical Neurophysiology*, *117*(7), 1574–1581. http://doi.org/10.1016/j.clinph.2006.03.011

Liu, C. C., Hosking, S. G., & Lenné, M. G. (2009). Predicting driver drowsiness using vehicle measures: recent insights and future challenges. *Journal of Safety Research*, *40*(4), 239–245. http://doi.org/10.1016/j.jsr.2009.04.005

Luo, X., Hu, R., & Fan, T. (2013). The driver fatigue monitoring system based on face recognition technology. In *2013 Fourth International Conference on Intelligent Control and Information Processing (ICICIP)* (pp. 384–388). http://doi.org/10.1109/ICICIP.2013.6568102

Revised Hours-of-Service Rule to Help Ensure Truck Drivers Get Adequate Rest. (2003, April 24). [Text]. Retrieved September 21, 2015, from http://www.fmcsa.dot.gov/newsroom/revised-hours-service-rule-help-ensure-truck-drivers-get-adequate-rest

Shahid, A., Wilkinson, K., Marcu, S., & Shapiro, C. M. (2011). Karolinska Sleepiness Scale (KSS). In A. Shahid, K. Wilkinson, S. Marcu, & C. M. Shapiro (Eds.), *STOP, THAT and One Hundred Other Sleep Scales* (pp. 209–210). Springer New York. Retrieved from http://link.springer.com/chapter/10.1007/978-1-4419-9893-4_47

Summary of Hours of Service Regulations. (2013, December 30). [Text]. Retrieved April 22, 2015, from http://www.fmcsa.dot.gov/regulations/hours-service/summary-hours-service-regulations

Zhang, M., Longhui, G., Wang, Z., Xu, X., Yao, B., & Zhou, L. (2014). Hybrid Model for Early Onset Prediction of Driver Fatigue with Observable Cues. *Mathematical Problems in Engineering*, *2014*, e385716. http://doi.org/10.1155/2014/385716

Chapter 5: Summary and Directions for Future Research

Summary and Directions for Future Research

5.0 Summary

This research has presented two different types of vehicle routing and scheduling planning problems: car shipping and fatigue-aware scheduling. In the car shipping variant, we explored the requirements to model and solve such a problem and introduced a novel backtracking restriction driven by our understanding of behavioral logistics. Further research into new types of backtracking restrictions and the underlying driver preferences is warranted.

In the fatigue-aware scheduling variant, we introduced the Truck Driver Scheduling Problem with Fatigue Monitoring (TDSPFM) by incorporating the Three Process Model of Alertness into the scheduling problem. We applied the TDSPFM to schedules to study the impact of different minimum alertness thresholds. Additionally, we quantified the effects of the driver's starting alertness level as it relates to their fatigue throughout the remainder of the work week. Further research to validate the TDSPFM using naturalistic driving data is warranted.

The ultimate goal of the TDSPFM is to develop safer schedules for trucker drivers. This is shared by government agencies who pass commercial vehicle driving regulations. However, the answer may not be more stringent regulations. It may be a data driven development and implementation of regulations. For instance, if one can measure how drowsy a driver is and predict when alertness may be an issue, then perhaps forcing mandatory rest periods at fixed intervals is not the safest or most efficient solution.

5.1 Future Research

In the car shipping problem, future research is warranted in the areas of problem size reduction and decision support system (DSS) integration. With respect to problem size, one of the challenges to solving a real-world problem car shipping problem is limiting the number of jobs considered, which directly affects the problem size. A reliable way to filter jobs that have little to no chance of being included in the solution should be developed. It is possible that other strategies for implementing the backtracking restriction could also help to reduce the problem size and thus other backtracking restriction techniques should be explored.

The multiple solutions produced by modifying the backtracking restriction parameter would serve a VRP DSS nicely. The development of or integration into an existing DSS should be explored. The techniques presented in chapter 2 would allow an independent trucker employed in the car shipping industry to find good routes and job selection using a laptop computer or tablet. A larger system, fronted by a web interface could solve many similar problems for a host of independent truckers using the same formulation and backed by cloud computing resources for a more complete DSS.

This research presents the first TDSPFM model in the literature, thus there is significant opportunity to validate and expand the model. This work shows how the TDSPFM performs on an existing real-world truck driver schedule extracted from an actual log book. We can further validate our model by using data sets being developed as part of existing research such as the "Commercial Motor Vehicle Driver Restart Study" sponsored by the Federal Motor Carrier Safety Administration. Additionally we could incorporate a study of driver schedules as part of a naturalistic driving study at the Virginia Tech Transportation Institute (VTTI).

There are two primary areas we can look at for expansion of the TDSPFM model: sleep variability and driving specific alertness modifications. We can look at sleep variability from both the standpoint of individual variations and specific rest period variations. The TPMA allows for user specific sleep functions to predict both recovery rates and sleep schedules. For instance, certain people may recover during sleep at different rates or be less alert after waking up at 3 a.m. to begin work. Additionally, our models assume non-stochastic sleep times during rest periods. However, practically speaking the amount of time that you actually sleep on a given night is dependent upon many factors (noise levels, comfort, temperature, etc.). We could incorporate a stochastic element into the model to represent how these factors affect sleep duration and recovery.

As a planning problem, we should look for areas where a decision support system (DSS) could aid in selecting safer schedules for drivers. This could allow for selection of parameter values such as a minimum acceptable alertness level, or it could alert schedulers to specific points along the route where alertness will be near an unsafe level. This could allow some specific action to be taken to reduce the risk of the driver being too fatigued. For example, the dispatcher could check in with the driver prior to the time where alertness is unsafe or extra precautions could be taken to increase the chances that the driver received a good night's rest the night prior. Other extensions to be considered include allowing for different alertness levels depending on the cargo being transported or the geographic areas through which the route passes.