Transportation Research Part E 116 (2018) 242–265

Contents lists available at ScienceDirect

Transportation Research Part E

journal homepage: www.elsevier.com/locate/tre

Optimizing for total costs in vehicle routing in urban areas

Jan Fabian Ehmke[a,*], Ann M. Campbell[b], Barrett W. Thomas[b]

[a] University of Magdeburg, Management Science Group, Universitätsplatz 2, 39106 Magdeburg, Germany
[b] The University of Iowa, Department of Management Sciences, United States

ARTICLE INFO

Keywords:
Emissions
Vehicle routing
Green logistics
Driver costs
Urban logistics

ABSTRACT

Minimizing cost is one of the most important objectives for logistics service providers, and it is not clear how an emphasis on minimizing emissions impacts costs. Most methodologies for routing currently minimize distance or travel time. This paper compares total cost (based on driver and fuel costs), fuel consumption/emissions, distance, and travel time for routes resulting from optimizing each of those measures. We explore the impact of multiple factors on these measures as well as the structure of the routes. Our results suggest that companies need rich cost models and routing algorithms with path flexibility to truly minimize total costs.

1. Introduction

With the increasing attention given to the impact that delivery vehicles have on the environment, many companies are seeking to manage their environmental impact while maintaining cost competitiveness. As examples, Walmart, General Mills, and Anheuser-Busch all have programs designed to reduce the emissions and thus fuel consumption of their delivery fleets (Hardcastle, 2015). Still, minimizing costs remains one of the most important objectives for logistics service providers, and it is not clear how minimizing emissions impacts total costs. Further, most methodologies for routing problems minimize neither cost nor emissions, but rather distance or travel time.

This paper examines a total cost objective that combines driver and fuel costs. We investigate the trade-offs that occur in total cost, time, distance, and fuel consumption (emissions) in urban areas when planning a priori routes for these different measures. In urban settings, vehicles must move at the speed of traffic and are also subject to the variability of those speeds. To understand the impact of this variability, we use a detailed fuel model and compare total cost minimized routes to those that optimize fuel consumption using a new simplified fuel consumption function. We experiment with different values for driver hourly cost, different values for fuel cost, customer geographies (inner city, suburban, mixed), customer load distributions (homogeneous versus heterogeneous), vehicle sizes (standard versus heavy), fleet compositions (homogeneous versus heterogeneous), traffic congestion (rush hour versus non-rush hour), and whether or not the vehicle is delivering or picking up loads at customers.

To derive our conclusions, we base our experiments on the road network of Stuttgart, Germany and a database of millions of speed observations. We solve instances on the network using an extension of the LANTIME routing heuristic (Maden et al., 2010), a well known heuristic for time-dependent vehicle routing problems. Our extension accounts for two features important to fuel consumption. First, the load of the vehicle changes as it visits customers, and the load impacts the fuel consumption. Second, there is variability in the speeds on arcs even at the same time of day. Because of the convexity of fuel-consumption curves with respect to speed, the failure to account for the speed variability by using an average speed value leads to an underestimation of the fuel usage. These two factors mean that the best path between two customers depends on the load and time of day of travel. Thus, in determining

* Corresponding author.
E-mail addresses: jan.ehmke@ovgu.de (J.F. Ehmke), ann-campbell@uiowa.edu (A.M. Campbell), barrett-thomas@uiowa.edu (B.W. Thomas).

https://doi.org/10.1016/j.tre.2018.06.008
Received 29 November 2017; Received in revised form 22 June 2018; Accepted 24 June 2018
1366-5545/

routes optimized for total cost or fuel consumption, we must solve for the best paths between customers on the routes. In the manner of Huang et al. (2017), we refer to the ability to change paths in response to these factors as path flexibility. To improve computational efficiency, we extend an earlier result that allows us to precompute and store a large number of these paths.

The most important conclusion from our work is that distance and time, objectives commonly used in routing optimization, are poor proxies for minimizing total cost. Minimizing distance can lead to significantly higher costs than minimizing total cost. Minimizing travel time can also lead to significantly higher fuel consumption than optimizing for total cost, especially with heterogeneous fleets. Also important for modern businesses facing pressure to be more sustainable, our results show that minimizing for total cost often increases fuel consumption only minimally over routes optimized for fuel consumption.

These conclusions suggest that companies who want to minimize costs need to focus on rich cost models that incorporate the cost of drivers' time as well as the cost of fuel consumption. It is also an important conclusion for routing research. Modeling total cost, particularly fuel consumption, introduces non-linearities not found in traditional routing objectives. As a result, there is a need for new routing methods that overcome the challenges.

Other contributions of this paper include:

– a literature review that captures the key differentiating features of the work done in the fast emerging area of emissions routing,
– an extension to our previous model for the vehicle routing problem to include fuel and driver costs in a total detailed cost objective function for urban areas,
– a new model for a fuel minimization objective function that is simplified from our previous work to help understand the value of our previous model,
– a tractable extension of the LANTIME heuristic that allows optimized paths between customers to be chosen that are both load-dependent and time-dependent (path flexibility), and
– an extensive set of experiments comparing the metrics achieved by optimizing for different objectives across many different test parameters using real world data.

The remainder of this paper is outlined as follows. In the next section, we review the relevant literature. Section 3 details our research questions and explains our methodology in detail. We introduce our experimental design and present results of our experiments in Section 5 and conclusions in Section 6.

2. Literature review

The vehicle routing problem is one of the most studied problems in operations research. However, until recently, most of the work focused on linear objectives, usually related to minimizing distance or travel time (Toth and Vigo, 2014). In recent years, researchers have begun to consider the impact of fuel consumption or emissions. Both fuel consumption and emissions are generally modeled as nonlinear functions of vehicle speed. In our paper, we model fuel consumption using the Comprehensive Emissions Model (CEM) described in Barth and Boriboonsomsin (2008). This model was originally designed to model the emissions and fuel consumption of heavy-goods vehicles and computes the fuel consumption on an arc as a function of speed, vehicle weight, and numerous vehicle and arc-specific constants. In this paper, we specifically implement the time-dependent version of the CEM found in Franceschetti et al. (2013). For comprehensive reviews of fuel consumption and emissions models, see Demir et al. (2011, 2014a).

The literature incorporating fuel consumption in routing and shortest path problems includes a wide variety of models and methodological approaches. Reviews can be found in Bektas et al. (2016), Demir et al. (2014a) and Lin et al. (2014). The work in our paper is most closely related to the work by Ehmke et al. (2016b) and Huang et al. (2017). Both papers require that vehicles travel at the speed of traffic. The need to travel at the speed of traffic comes from the papers' focus on urban settings. On most roads in urban settings, not only are the cost- and fuel-optimal speeds above the legal speed limit, but congestion forces vehicles to travel at a given speed. In these settings, both Hwang and Ouyang (2015) and Ehmke et al. (2016a) demonstrate that accounting for the variability in travel speeds leads to different path choices than when solving deterministic equivalents.

Important to this work is the path-finding methods introduced in Ehmke et al. (2016a). Ehmke et al. (2016a) introduces two heuristic methods for finding expected shortest-emission paths when the travel speed distributions are time dependent. The paper introduces two methods for finding minimum fuel consumption minimizing paths. The first method incorporates sampling into an $A^{\star}$-based algorithm. The second uses an averaging technique that captures the impact of speed variability on the fuel consumption on an arc. Results in Ehmke et al. (2016a) demonstrate that the information loss resulting from the modelling assumptions of the latter method lead to a minimal difference in solution quality. The solution methodology for the COST and FUEL objectives in our paper takes advantage of the averaging method. We note that Wen et al. (2014) consider the case of finding minimum cost paths in a time-dependent network with congestion charges. However, unlike Ehmke et al. (2016a), Wen et al. (2014) do not consider variation in speeds nor the nonlinearities introduced by minimizing fuel consumption. This paper incorporates the time-dependent path-finding first introduced in Ehmke et al. (2016a).

Ehmke et al. (2016b) explore the impact of multiple pickups and vehicle load on expected emissions-minimized routes. Ehmke et al. (2016b) recognizes that, when considering fuel consumption in the objective, the best paths between customers change as a function of the time-dependent distribution of the speeds as well as the load of the vehicle at the time that the path is traversed. This feature is known as path flexibility. Ehmke et al. (2016b) use a generalization of the path-finding methods of Ehmke et al. (2016a) to determine the appropriate emissions-minimizing paths between each pair of customers given the current time and load of the vehicle. To reduce the computational burden of computing these paths at runtime, Ehmke et al. (2016b) introduce a theoretical result that