European Journal of Operational Research 000 (2018) 1–16

Contents lists available at ScienceDirect

European Journal of Operational Research

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

Production, Manufacturing and Logistics

The Dynamic Dispatch Waves Problem for same-day delivery

Mathias A. Klapp[a,*], Alan L. Erera[b], Alejandro Toriello[b]

[a] Engineering School, Pontificia Universidad Católica de Chile, Santiago, Chile
[b] H. Milton Stewart School of Industrial and Systems Engineering, Georgia Institute of Technology, Atlanta, GA 30332, USA

ARTICLE INFO

Article history:
Received 29 May 2017
Accepted 17 May 2018
Available online xxx

Keywords:
Logistics
Same-day delivery
Vehicle routing
Approximate dynamic programming

ABSTRACT

We study same-day delivery (SDD) systems by formulating the Dynamic Dispatch Waves Problem (DDWP). The DDWP models an order dispatching problem faced by a distribution center, where orders arise dynamically throughout a service day and must be delivered by day's end. At each decision epoch (wave), the system's operator chooses whether or not to dispatch a single vehicle loaded with orders ready for service in order to minimize vehicle travel costs and penalties for unserved requests. We formulate an arc-based integer programming model and design local search heuristics to solve a deterministic DDWP where order arrival times are known in advance. We use the deterministic variant to design an *a priori* solution approach, and provide two approaches to obtain dynamic policies using the *a priori* solution. We test and compare solution approaches on two sets of instances with different geography scenarios, size, information dynamism, and order timing variability. The computational results suggest that our best dynamic policy can reduce the average cost of an *a priori* policy by 9.1% and substantially improves the fraction of orders delivered (order coverage), demonstrating the importance of reactive optimization for dynamic SDD services. We also analyze the tradeoff between two common SDD objectives: total cost minimization versus order coverage maximization. We find structural differences in the dispatch frequency and route duration of solutions for the two different objectives, and demonstrate empirically that small increases in order coverage may require substantial increases in vehicle travel cost.

© 2018 Elsevier B.V. All rights reserved.

1. Introduction

Same-day delivery (SDD) is increasingly being offered by retailers and logistics service providers to expand e-commerce. The internet sector represented over 7% of the U.S. retail industry and was expected to grow 14% annually in 2015 (DeNale, Liu, & Weidenhamer, 2015; Lindner, 2015). We define SDD as a distribution service that prepares, dispatches and delivers orders to the customer's location on the same day the customer places the order. Amazon, a provider of SDD, has implemented the service in more than 25 U.S. metropolitan areas as of October 2016, and has also piloted "Prime Now," an even faster one-hour delivery service. These services are designed to satisfy demand for instant gratification when ordering consumer products and at the same time to discourage physical store visits (Klapp, Erera, & Toriello, 2018; Voccia, Campbell, & Thomas, 2015). Like all urban delivery operations, same-day delivery operations can be costly due to small order sizes to be delivered to many geographically dispersed locations, and they are additionally challenged by having less time to consolidate orders into effective vehicle routes.

A core logistics process within any delivery system is order distribution from stocking locations to delivery locations. Distribution decisions can be divided into dispatch decisions, which select the dispatch times and the orders to be delivered in each delivery vehicle trip, and routing decisions, which sequence the delivery locations for each dispatch. Dispatch time selection for same-day services is often more challenging than traditional distribution, since vehicles may not be dispatched simultaneously at the beginning of the operating day and may also be reused during the day. It may also be sensible to not dispatch a ready order if there is reasonable likelihood that later (unknown) orders may be consolidated with it. In this research, we explicitly consider these challenges. We consider problems where dispatch times are selected from a fixed and finite number of candidates, both for simplicity and to maximize efficiencies gained by organizing warehouse activity using *wave picking*. In wave picking, the workload of picking activities is organized and balanced by pick wave, *i.e.*, a "scheduled chunk of work" (Bartholdi & Hackman, 2008). Pick waves are desirable for many reasons, including constraints associated with shift start times, workday length limits, operation times within the warehouse, and other restrictions related to the

---

[*] Corresponding author.
E-mail addresses: maklapp@ing.puc.cl (M.A. Klapp), aerera@isye.gatech.edu (A.L. Erera), atoriello@isye.gatech.edu (A. Toriello).

https://doi.org/10.1016/j.ejor.2018.05.032
0377-2217/© 2018 Elsevier B.V. All rights reserved.

Please cite this article as: M.A. Klapp et al., The Dynamic Dispatch Waves Problem for same-day delivery, European Journal of Operational Research (2018), https://doi.org/10.1016/j.ejor.2018.05.032

organization of work. Typically, each wave is designed to last between one and four hours (Piasecki, 2009); furthermore, in some consolidation centers, it is common to dispatch a vehicle between two and three times per day (van Heeswijk, Mes, & Schutten, 2017; Nesterova & Quak, 2015).

The Dynamic Dispatch Waves Problem (DDWP) is an order delivery problem with dynamic dispatch and routing decisions for a single vehicle during a fixed-duration operating period (*i.e.*, a day) partitioned in *W dispatch waves*. For our purposes, a dispatch wave is analogous to a pick wave in the context of vehicle dispatch decisions. The start of each wave is a time point when picking and packing of a set of orders is completed, and a vehicle (if available at the depot) can be loaded and dispatched from the depot to serve a subset of open delivery orders, or wait for the next wave. For simplicity, we refer to the start time of a dispatch wave, *i.e.*, the potential decision epoch, as a “wave” as well. Open orders are defined as those ready to be dispatched and not previously served. After the vehicle completes a dispatch route, it returns to the depot and is ready to be dispatched again. At each wave, complete information is known for all open orders, and probabilistic information is available describing potential future orders. The objective is to minimize expected vehicle operating costs and penalties for open orders that remain unserved at the end of the operating period.

In the DDWP, we assume unserved realized orders are not formally rejected until the end of the day. Compared to a framework where an order’s acceptance decision must be executed immediately, our setting gives the service provider an extra degree of operational flexibility, and models a system in which all realized orders are ultimately served. That is, at the end of the planning period, all rejections are outsourced to a third-party logistics provider at an order-dependent cost. In Klapp (2016), we study a model for the case where an order must be immediately accepted or rejected upon realizing.

The DDWP captures two fundamental and important tradeoffs in same-day distribution. First, there is a tradeoff between waiting and dispatching a vehicle at the start of each wave. When a vehicle is dispatched, the set of open orders is reduced, but the opportunity to observe and serve future orders near ones in the current route is lost. Conversely, when the vehicle is not dispatched, the time remaining to serve the set of open and future orders is reduced. Second, there is a tradeoff between dispatching a long route that serves many orders versus a short one with fewer orders. The former consumes more time and keeps the vehicle away from the depot longer, but requires less time per customer visited due to routing economies. A shorter route with fewer orders requires more time per customer, but enables the vehicle to be reused sooner.

1.1. Contributions

In Klapp et al. (2018), we defined the one-dimensional DDWP for a simplified SDD distribution system with delivery locations on the one-dimensional line. In this paper, we formulate the single vehicle DDWP in a general network topology, and provide a more realistic model of same-day delivery operations in a typical road network. We capture fundamental tradeoffs in dynamic dispatch decision-making, and model a canonical prize-collecting version of the DDWP. We consider the following to be our primary contributions.

1. We formulate an IP model for a deterministic variant where all order arrival times are known in advance, which we leverage to provide *lower bounds* for the stochastic-dynamic case via information relaxation and simulation.
2. We use the deterministic model to find an *optimal a priori solution* to the stochastic variant, by showing that the *a priori* optimization problem is equivalent to an instance of the deterministic variant with an expanded customer set and adjusted penalties. To our knowledge, this is the first time a simulation-free optimal *a priori* policy is presented for a model that includes route sequencing decisions in the same-day delivery distribution literature. We also design construction and local search heuristics based on the prize-collecting Traveling Salesman Problem to complement commercial IP solvers and speed up the identification of solutions to problem instances.
3. We develop two approaches to obtain *dynamic policies* using the *a priori* model. The first uses a rollout scheme to dispatch according to an updated *a priori* solution, and the latter relies instead on fast heuristic modifications to the initial *a priori* solution.
4. We prove that the expected cost of the rollout policy is no higher than the optimal *a priori* policy cost. We later confirm the benefits of dynamic policies with computational experiments. Experiments show that the percentage reductions in optimality gap and operating cost from an *a priori* policy to a dynamic rollout policy average 47.6% and 9.1%, respectively, while the order fill rate improves by 4.6%. The marginal benefits of dynamic policies increase for instances with greater order arrival variability and less information disclosed before the start of the operation. Performance measures for dynamic policies also exhibit less variability across instances.
5. Finally, we analyze the tradeoff between two common objectives in SDD, minimizing total costs (including vehicle travel time) versus maximizing order coverage, in an empirical study. One might suppose that these two objectives lead to similar solutions, since low-cost routes create more available vehicle time to serve additional orders. In contrast, we instead find fundamentally different solution structures for the two cases in terms of the number of vehicle routes dispatched, route length and initial wait at the depot. Results indicate that one should expect significant sacrifices in vehicle routing efficiency in order to maximize the order fill rate, and the cost of an additional customer covered becomes more expensive as order coverage increases.

The remainder of the paper is organized as follows. A literature review is presented in Section 1.2. Section 2 defines notation and formulates the DDWP, Section 3 covers the deterministic problem, and Sections 4 and 5 respectively cover *a priori* and dynamic policies. Finally, Section 6 outlines the results of a computational study, and we conclude with Section 8.

1.2. Literature review

The DDWP can be classified within the broad family of stochastic vehicle routing problems (VRP) that are extensions of the deterministic VRP (Golden, Raghavan, & Wasil, 2008; Toth & Vigo, 2014). The simplest models for this problem family are *a priori* optimization models (Campbell & Thomas, 2008a; Cordeau, Laporte, Savelsbergh, & Vigo, 2006; Gendreau, Laporte, & Séguin, 1996; Jaillet, 1988), where an initial solution is designed before the operation starts and then pre-defined recourse rules, *i.e.*, simple plan corrections, are used to (potentially) modify it during operation as information is disclosed. More complex models for dynamic-stochastic VRPs are dynamic policies that adapt to revealed information during the operating period, and allow for re-optimization of structural routing and scheduling decisions; see Bent and van Hentenryck (2004), Larsen, Madsen, and Solomon (2008), Pillac, Gendreau, Guéret, and Medaglia (2013), and Thomas (2010). A restricted case of dynamic policies are *a priori* rollout policies that proved

Please cite this article as: M.A. Klapp et al., The Dynamic Dispatch Waves Problem for same-day delivery, European Journal of Operational Research (2018), https://doi.org/10.1016/j.ejor.2018.05.032

Download English Version:

https://daneshyari.com/en/article/6894451

Download Persian Version:

https://daneshyari.com/article/6894451

Daneshyari.com