A Multi-Period Vehicle Routing Problem with State-Dependent Service Times for Last-Mile Distribution Networks

Christian Truden[1], Margaretha Gansterer[1], Dominik Loske[2], and Matthias Klumpp[3,4]

[1]Department of Economics, Analytics, and Operations Research, University of Klagenfurt, Klagenfurt, Austria ✉ christian.truden@aau.at, margaretha.gansterer@aau.at
[2]Institute of Production and Supply Chain Management, Technical University of Darmstadt, Darmstadt, Germany, ✉ loske@pscm.tu-darmstadt.de
[3]HumanTech CORE Lab, School of Management, Politecnico di Milan, Milano, Italy, ✉ matthias.klumpp@polimi.it
[4]Retail Logistics & Innovation Lab, Department of Business Studies and Economics, University of Bremen, Bremen, Germany

1 Introduction

Vehicle routing problems (VRPs) form a central class of combinatorial optimization problems in logistics and transportation. Standard formulations typically assume that service times at customer locations are deterministic and independent of past decisions. In practice, however, service times often exhibit systematic heterogeneity driven by accumulated familiarity with delivery locations, unloading procedures, and interpersonal coordination between delivery drivers and customers. Such effects can account for a substantial variant among of the expected service times. While uncertainty in demand and travel times has been extensively studied, comparatively little attention has been paid to endogenous, history-dependent service times in vehicle routing models. We address this gap by introducing a multi-period vehicle routing problem in which service times depend on a familiarity state that evolves endogenously as a function of past routing decisions, i.e., state-dependent cost.

2 Multi-Period Vehicle Routing Model

We consider a planning horizon consisting of discrete periods $p = 1, \ldots, P$. In each period, a set of customer locations I must be served from a central depot using a fleet of vehicles V operated by drivers D. Customer demands, time windows, and vehicle availability may vary by period, while the transportation network remains fixed.

To model history-dependent service times, we introduce a familiarity matrix $(\lambda^p)_{di} \in [0,1], \quad d \in D,\ i \in I$, where λ^p_{di} represents the degree of familiarity of driver d with customer location i in period p. A value of $\lambda^p_{di} = 0$ indicates no familiarity, while $\lambda^p_{di} = 1$ denotes full familiarity.

Service times are assumed constant within each period but depend on the familiarity state. Let k^p_i denote the commodity size to be delivered to customer i in period p. The service time incurred when driver d serves customer i using vehicle v is given by $s^p_{idv} = s_{idv}(\lambda^p_{di}, k^p_i) \in \mathbb{R}^+$. In the computational study, we adopt the following parametric specification:

$$s_{idv} = k^p_i \cdot t^{\text{palette}} + (\underline{t}^{\text{arrive}}_i + \underline{t}^{\text{depart}}_i) + (1 - \lambda_{di}) \cdot (\widetilde{t}^{\text{arrive}}_i + \widetilde{t}^{\text{depart}}_i) + \lambda_{di} \cdot t^{\text{social}} + m_v, \tag{1}$$

where arrival and departure times are decomposed into lower bounds and variable components, t^{social} captures social interaction time, and m_v penalizes maneuvering effort for larger vehicles.

For a fixed familiarity matrix λ^p, each period induces a classical vehicle routing problem with time windows and deterministic service times s^p_{idv}. Let $\theta^{*p}(V, D, K_p)$ denote an optimal routing solution in period p, where K_p represents the period-specific demand profile.

After each period, the familiarity matrix is updated based on the realized routing solution as $\lambda^{p+1} = \lambda^p[\lambda^p, \theta^{*p}(V, D, K_p)]$. We assume a linear learning-and-forgetting mechanism in which familiarity increases when a driver serves a location and decays otherwise, i.e., $\lambda^{p+1}_{di} := \lambda^p_{di} + \Delta^+ \mathbb{I}_{di}(\theta^{*p}) + \Delta^- \bar{\mathbb{I}}_{di}(\theta^{*p})$, where

$\mathbb{I}_{di}(\theta^{*p})$ equals one if driver d visits location i in period p, and zero otherwise. Parameters $\Delta^{+} > 0$ and $\Delta^{-} < 0$ control learning and forgetting rates. The resulting formulation constitutes a multi-period vehicle routing problem in which routing decisions affect future service times through the familiarity state, while preserving a classical VRPTW structure in each period.

3 Computational Study

We conduct a computational study to elaborate the structural implications of state-dependent service times in a multi-period routing context. The study is based on empirical data from a grocery distribution network serving stores within a 30 km radius of a central depot. Deliveries are performed six days per week using a heterogeneous fleet of four vehicle types. To capture systematic variation in demand and routing conditions, we construct six representative instances corresponding to weekdays from Monday to Saturday. The update parameters Δ^{+} and Δ^{-} are calibrated using historical delivery records provided by the industry partner. Due to the linear structure of the update rule, calibration is equivalent to a regression model implemented as a single-layer feed-forward neural network. The model is simulated over 60 periods, corresponding to ten weeks of operation.

Initial analysis shows that approximately 80 % of total route duration is attributable to service times, underscoring the relevance of modeling service-time heterogeneity. We compare the proposed location-based update rule with an unbiased baseline in which familiarity values are evenly distributed across all driver–location pairs. Figure 1 reports relative savings in the variable component of service times across

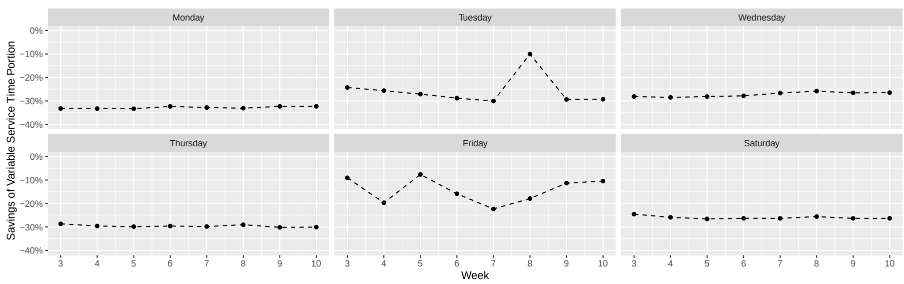

Figure 1: Relative savings from modeling the variable component of service times across weekdays.

weekdays. Savings exceed 33 % on Mondays and range between 15 % and 25 % on other weekdays, with Fridays exhibiting the largest variability. These effects are stable across weeks, indicating that routing decisions induce persistent changes in node service costs through the familiarity-driven state update mechanism. Aggregated over the planning horizon, the average reduction in total service time amounts to 5.5 %.

4 Conclusion

We introduce a multi-period vehicle routing problem with endogenous, state-dependent service times. By modeling service times as a function of a familiarity state between drivers and delivery nodes that evolves according to past routing decisions, the proposed framework captures learning and forgetting effects within a multi-period routing context. The formulation preserves classical VRPTW structure at the single-period level while introducing intertemporal coupling through a simple and flexible state update rule.

The assumed state-dependent service time model opens several avenues for future research. Moving beyond the current myopic updating of service times, anticipatory planning approaches that explicitly incorporate learning, along with stochastic modeling and sustainability considerations such as synchronized electric vehicle recharging, could further improve the robustness and efficiency of distribution operations.