TOP
https://doi.org/10.1007/s11750-023-00656-6

ORIGINAL PAPER

Reoptimisation strategies for dynamic vehicle routing problems with proximity-dependent nodes

Tiria Andersen[1] · Shaun Belward[1] · Mangalam Sankupellay[1] · Trina Myers[2] · Carla Chen[1]

Received: 1 September 2022 / Accepted: 15 March 2023

Abstract

Autonomous vehicles create new opportunities as well as new challenges to dynamic vehicle routing. The introduction of autonomous vehicles as information-collecting agents results in scenarios, where dynamic nodes are found by proximity. This paper presents a novel dynamic vehicle-routing problem variant with proximity-dependent nodes. Here, we introduced a novel variable, *detectability*, which determines whether a proximal dynamic node will be detected, based on the sight radius of the vehicle. The problem considered is motivated by autonomous weed-spraying vehicles in large agricultural operations. This work is generalisable to many other autonomous vehicle applications. The first step to crafting a solution approach for the problem is to decide *when* reoptimisation should be triggered. Two reoptimisation trigger strategies are considered—exogenous and endogenous. Computational experiments compared the strategies for both the classical dynamic vehicle routing problem as well as the introduced variant. Experiments used extensive standardised vehicle-routing problem benchmarks with varying degrees of dynamism and geographical node distributions. The results showed that for both the classical problem and the novel variant, an endogenous trigger strategy is better in most cases, while an exogenous trigger strategy is only suitable when both detectability and dynamism are low. Furthermore, the optimal level of detectability was shown to be dependent on the combination of trigger, degree of dynamism, and geographical node distribution, meaning practitioners may determine the required detectability based on the attributes of their specific problem.

Keywords Dynamic · Vehicle routing · Proximity-dependent · Autonomous vehicles

Shaun Belward, Mangalam Sankupellay and Trina Myers contributed equally to this work.

✉ Tiria Andersen
tiria.andersen@my.jcu.edu.au

✉ Carla Chen
carla.ewels@jcu.edu.au

Extended author information available on the last page of the article

Published online: 10 May 2023

Mathematics Subject Classification 90-05 · 90-08 · 90-10

1 Introduction

Autonomous vehicles are bringing new modes of transport into the transportation world (Savelsbergh and Van Woensel 2016). When planning routes for autonomous vehicles traveling to multiple locations or covering a large geographical area, minimising the vehicle's travel distance is essential to keeping operational costs low. As such, there is a growing body of research in Vehicle Routing Problems (VRP) for autonomous vehicles (e.g., Poikonen et al. 2017; Boysen et al. 2018; Hyland and Mahmassani July 2017; Bsaybes et al. 2019; Chen et al. 2021).

Due to the establishment of fast information transfer within the last two decades, dynamic VRPs (DVRPs) are becoming the norm (Gendreau et al. 1999; Savelsbergh and Van Woensel 2016). Traditionally, DVRPs consist of human-driven vehicles visiting dynamically appearing customer locations. DVRPs with autonomous vehicles, on the other hand, do not need to accommodate human drivers. Autonomous vehicles can, therefore, travel over previously inaccessible terrain—be it air, land or sea. The diversification of application contexts mean that customers (or *nodes*) can be anything from buildings, vehicles, injured disaster victims, weeds, wildlife, to locations of interest in general.

This study is motivated by a real life application of autonomous weed-spraying vehicles used to survey agricultural land and address weed infestations in North Queensland, Australia. Australia has some of the largest cattle stations in the world, with individual properties spanning thousands of square kilometers (for example, Australia's largest cattle station is larger than Israel). These agricultural properties are often infested by large invasive plant species, which impedes effective operations, and are considered an issue of importance due to their significant effect on the agricultural industry. Curtailing the spread of invasive species is an ongoing effort by both farmers and government programs. Due to the clustered nature of weed growth, spot-spraying herbicide manually from an all-terrain vehicle is an efficient method that minimises herbicide use. However, the task can be labour intensive and hazardous due to the size of properties and continuous contact with chemicals.

For this reason, prototype autonomous weed-spraying vehicles are emerging in the agricultural market. These autonomous vehicles are often equipped with camera vision, and are capable of identifying new weeds visible within a certain proximity. In contrast to traditional DVRPs which relied on customers dynamically providing location information, vehicles are now information-collecting agents dealing with nodes that are proximity-dependent. As such, this paper introduces the DVRP with *proximity-dependent nodes*.

The DVRP with proximity-dependent nodes can be generalisable to many other applications. Here are some examples where the exact location of nodes are unknown but detected by the proximity to bots. For example, autonomous swarms sent out to locate and dispense first aid to injured victims after a natural disaster; small ground-based bots to search and inspect potential faults in industrial construction and maintenance; cleaner micro-bots sterilising high-contact hotspots and addressing irregular spillage; urban drones patrolling and addressing events of concern as a first-responder in highly populated cities; conservation subs tracking and tagging marine life or endangered species; survey rovers exploring remote locations, taking measurements and investigating points of interest. Solving the DVRP with proximity nodes is thus important for building towards operating large-scale fleets of adaptable autonomous vehicles that can address known points of interest while flexibly exploring and responding to their environment.

When solving a DVRP, the first step is to characterise the problem attributes, and determine whether they are controllable. Depending on the context, problem attributes are either uncontrollable, and thus part of the informational process, or controllable and thus optimised as part of the decisional process (Gendreau et al. 2016).

The informational process encapsulates how, where, and when information is revealed or obtained through time. The attributes of the informational process essentially determine the dynamism of DVRP. The most significant attribute that determines dynamism is the proportion of dynamic nodes—*degree of dynamism* (dod)—which was introduced by Lund (Lund et al. 1996), and further extended by Larsen(Larsen 2000; Larsen et al. 2007). Another relevant dynamic attribute is the geographical distribution of nodes (Pillac et al. 2013). Because the detection of dynamic nodes are based on proximity, the amount of dynamic nodes detected is directly affected by the geographical distribution of nodes. The geographical node distribution is generally sorted into three categories: uniform random (R), random-clustered (RC) and clustered (C) (Uchoa et al. 2017). The clustered node distribution is especially relevant in this study, as it reflects the biological reproduction mechanisms of weeds.

The decisional process is the controllable response to revealed information. The first step in the decisional process is the trigger, which determines *when* routes are reoptimised (Ritzinger et al. 2016). Two trigger types are considered in this study: *endogenous* triggers and *exogenous* triggers. An endogenous reoptimisation is triggered by internal events, such as when a vehicle reaches a node. As such, once the destination node is fixed, the vehicle route cannot be reoptimised or changed until the vehicle reaches the destination node. An exogenous reoptimisation is triggered by external events, such as when a dynamic node appears. As such, the vehicle is allowed to immediately divert to a different destination in response to the new information.

Using an exogenous trigger means the vehicle is flexible to new information, but may result in premature decision-making. On the other hand, an endogenous trigger can exploit the information accumulated as the vehicle travels to its fixed destination, but may not react immediately when necessary. Studies suggest that trigger performance depends on a problem's attributes (Regan et al. 1998; Moretti Branchini et al. 2009; Ulmer et al. 2017). Specifically, choosing the appropriate trigger based on the geographical node distribution and dod can result in large improvements in optimisation (Ulmer et al. 2017).

This study aims to compare exogenous and endogenous trigger suitability for the DVRP with proximity-dependent nodes. To model the degree of proximity, *detectability* is introduced as an new attribute. Detectability is defined as a radius of sight, within which a proximal dynamic node will be detected. Trigger performance will be compared on varying dod, geographical node distribution, and detectability, for both the traditional DVRP and the DVRP with proximity-dependent nodes. Trigger performance is assessed by two criteria: solution distance, and the proportion of dynamic nodes detected.

Computational experiments are tested on the extended Set X benchmarks by Uchoa et al. (2017) to ensure an extensive variation in geographical distribution. The contributions of this paper can be summarised as follows:

1. The DVRP with proximity-dependent nodes is introduced, and its associated DVRP attribute *detectability* is defined.
2. The performance of endogenous and exogenous triggers are compared across dod, geographical node distribution, and detectability. Computational experiments are run on the well-known Uchoa benchmarks for both the basic DVRP and the introduced variant.
3. The relationship between detectability and dod as well as geographical node distribution is analysed, helping practitioners adjust detectability effectively to suit the characteristics of their application problem.

The remainder of this paper is outlined as follows: Sect. 2 gives a brief summary of the background literature related to this study; Sect. 3 defines the DVRP with proximity-dependent nodes. Section 4 describes the computational experiments and analysis techniques used. The computational results are presented in Sect. 5, followed by a discussion in Sect. 6. Finally, Sect. 7 gives the concluding remarks.

2 Related works

While proximity-dependent dynamic nodes are potentially novel, the concept of vehicles collecting information itself is not new. Psaraftis (1995) introduced information availability as a dynamic attribute that is either *local* or *global*. Global information availability means that dynamic information is available globally to the system, independent of the system state: e.g., a customer calling the central dispatcher to request a delivery. Most DVRP literature deals with global information availability (Ferrucci 2013). Local information availability is where dynamic information is available depending on the system state: e.g., when the vehicle is in near proximity to the source of information. Local information availability is common in case of the DVRP with stochastic travel times, where travel time is only revealed when the vehicle reaches an area with traffic congestion (Van Woensel et al. 2008). Similarly, in the DVRP with stochastic demand, node locations are known a priori, but the *demand* required by the node is only revealed when the demand-delivering vehicle arrives at the node location (Dror et al. 1989). In contrast, there is a significant lack of literature on DVRP with dynamic nodes under local information availability. Other variants of similar nature include the information-collecting VRP (Al-Kanj et al. 2016) and the covering tour problem (Margolis et al. 2022). However, proximity-dependent nodes are distinctive in the sense that the detection of dynamic nodes is purely reliant on the vehicle's actions, a mechanism not seen in the literature Psaraftis et al. (2016); Braekers et al. (2016).

Perhaps this is due to the historical use of human vehicle drivers. The visual and cognitive distraction of detecting nodes, registering information through a human-to-computer interface, and diverging from current routes, all reduces the safety and efficiency of the transportation processes (Ferrucci and Bock 2015). Autonomously driven vehicles do not have this concern. Autonomous vehicles are uniquely equipped to find dynamic nodes; with the assistance of computer vision, their onboard computers are capable of recognising, registering, and reoptimising on the go. Naturally, systems with a mix of both local and global information sources are to be expected in future. For now, the study will address locally collected information in isolation.

When considering the body of literature surrounding reoptimisation triggers, comparison of triggers is mostly found in *DVRPs with Time Windows* (DVRPTW). This is because DVRPTWs prioritise visiting customers within an urgent time window, so the timing of reoptimisation is crucial.

Studies with time windows often use DVRPTW benchmarks by Solomon (1987). For example, the exogenous trigger was reported by two studies as superior to the endogenous trigger, across all geographical distributions when compared on Solomon's DVRPTW benchmarks (Ichoua et al. 2000; Lorini et al. 2011). Both studies minimised lateness to serve customers and distance travelled as their objective functions. Note that the studies still differed in some aspects. For Ichoua et al. (2000), the trigger was indirectly exogenous; a reoptimisation was triggered at a dynamic interval calculated from the frequency of dynamic nodes appearing so far. Lorini et al. (2011) incorporated dynamic travel times in addition to dynamic nodes. Interestingly, they found that the exogenous trigger specifically excelled on randomly distributed (R) instances.

However, other studies on different benchmarks found that using the exogenous trigger only yielded improvements under certain conditions. Regan et al. (1998) and Ulmer et al. (2017) reported improvements when using the exogenous trigger, but only for instances with R distributions at low dod. Ulmer et al. (2017) compared triggers on their own DVRPTW benchmarks, and performance measured was by the proportion of customers served. Regan et al. (1998) tested the triggers on different sized vehicle fleets for the *DVRP with Pickup and Delivery with Time Windows* (DVRPPDTW). Trigger performance was measured by revenue proportional to distance traveled, minus a penalty for slack time when vehicles were idle. In their study, the endogenous trigger was better on instances with a small-medium fleet of vehicles at medium dod, while the exogenous trigger was superior on instances with a large fleet of vehicles at low dod.

In contrast, Moretti Branchini et al. (2009) found that using an exogenous trigger yielded insignificant improvements across the board, and especially for instances with randomly clustered nodes. The study's objective function minimised both distance traveled and slack time when vehicles were idle.

In general, the exogenous trigger seems to perform better when the geographical distribution of nodes are more random, and the dod is low. However, studies differ on whether using an exogenous trigger is significantly better than an endogenous trigger. The lack of consensus may be due to differing problem variants, objective functions, and benchmark instances.

There is a distinct gap in the literature on comparing endogenous and exogenous triggers on a standardised DVRP benchmark. To address this gap, this study compares the performance of exogenous and endogenous triggers on the extended Set X benchmarks by Uchoa et al. (2017), which contains an extensive set of realistic instances for each geographical distribution. The effects of dod and geographical distribution on trigger performance for both the traditional global DVRP as well as local DVRP with proximity-dependent dynamic nodes were investigated.

3 Problem definition

The classical VRP can be defined as a weighted graph $G = (V, E)$ where *nodes* is represented by the set of vertices $V = \{1, 2, 3 \dots n\}$. Paths between ith and jth customers are given by edge e_{ij} in edge set E and each edge has an associated *travel cost* c_{ij} and *travel time* t_{ij}. A vehicle, starting at the *depot* (typically node/vertex 0), is to service all nodes exactly once. Thus the objective is to create a feasible vehicle routing *solution* S, which sequentially addresses all nodes at minimum cost.

In a Capacitated VRP, each ith node has an associated *demand* q_i, and the vehicle has a limited *demand capacity* Q. The demand is sourced from the depot, and if the vehicle's demand capacity Q is smaller than the total demand of all nodes, then several routes will be needed.

The dynamic extension of the CVRP is the Dynamic VRP (DVRP), where some nodes are known a priori, while others are discovered dynamically as the vehicle is executing its route. To serve dynamic nodes enroute, it is assumed that the demand is homogenous (the same product/service for every node, see Eglese and Zambirinis 2018). The partition between a priori nodes and dynamic nodes is based on the degree of dynamism (dod). When the dod is zero, there are no dynamic nodes and the problem is static. When the dod is one, the problem is fully dynamic, or *online*, such that none of the nodes are known beforehand. As the dod increases, the proportion of dynamic nodes increases, and vehicle routes inevitably become longer and less efficient due to having less a-priori information.

The overall progression of node status during the simulation is as follows. Let V be the set of all nodes, partitioned into two subsets based on the dod: known nodes $V_{\text{a priori}}$ and dynamic nodes V_{dynamic}. At the beginning of the simulation, the static problem is solved and all known nodes in $V_{\text{a priori}}$ are assigned to routes, and moved to N_{assigned}. (Route feasibility is assumed).

As the vehicle executes its routes and comes within proximity of dynamic nodes, the detected dynamic nodes are moved from V_{dynamic} and cached in N_{dynamic}. When a reoptimisation is triggered, the cached nodes in N_{dynamic} will be assigned and thus

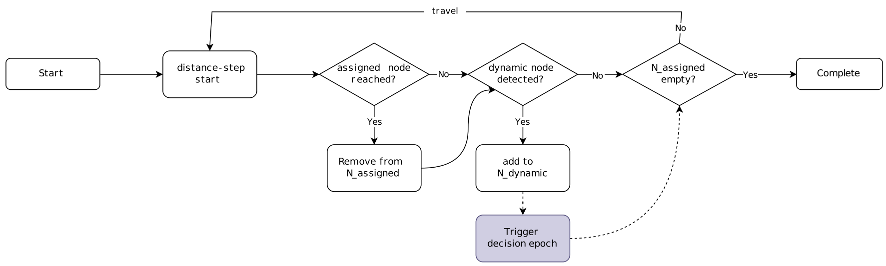

Fig. 1 Simulation process with exogenous trigger. Each distance step checks if an assigned node is reached first, then checks to detect dynamic nodes. A decision epoch is triggered if a dynamic node is detected. If multiple dynamic nodes are detected, all will be added to N_{dynamic} before triggering the decision epoch. The decision epoch empties N_{dynamic} and reoptimises routes to include dynamic nodes. (Parts unique to the exogenous trigger are denoted in dotted lines)

moved into N_{assigned}. Finally, when an assigned node is visited by the vehicle, it is moved from N_{assigned} to N_{visited}.

Because dynamic nodes are detected based on proximity to the vehicle rather than time, the progression of the vehicle was simulated as distance-steps rather than time-step: $\sum_i d_i = D$, where d is the distance travelled at distance-step i, and D is the total distance travelled. Thus, at the beginning of the simulation $d = d_0$, $N_{\text{assigned}} = V_{\text{a priori}}$, $N_{\text{dynamic}} = \emptyset$, and $N_{\text{visited}} = \emptyset$.

3.1 Detecting dynamic nodes

At the start of its route, the vehicle coordinates (x, y) will be located at the depot. Once the vehicle leaves the depot, it can detect dynamic nodes through local information availability.

To parameterise the local information availability, *detectability* was introduced as a dynamic attribute. The larger the detectability, the more dynamic nodes will be detected simultaneously as the vehicle moves. If detectability is small, only dynamic nodes directly next to the vehicle's path will be detected, and there is no guarantee all dynamic nodes will be detected. If detectability is larger than the problem area, the entire problem area will be in the vehicle's area of sight; all dynamic nodes will be detected immediately, rendering the problem static.

For this problem, the area of sight is given as a circle around the vehicle with radius r. At each distance step, the distance between all dynamic nodes in V_{dynamic} and (x, y) are calculated, and if the distance is less than r, the dynamic node is detected. If detected, the dynamic nodes are removed from V_{dynamic} and cached in N_{dynamic}. At the next decision epoch, N_{dynamic} will be emptied and the dynamic nodes assigned into a route.

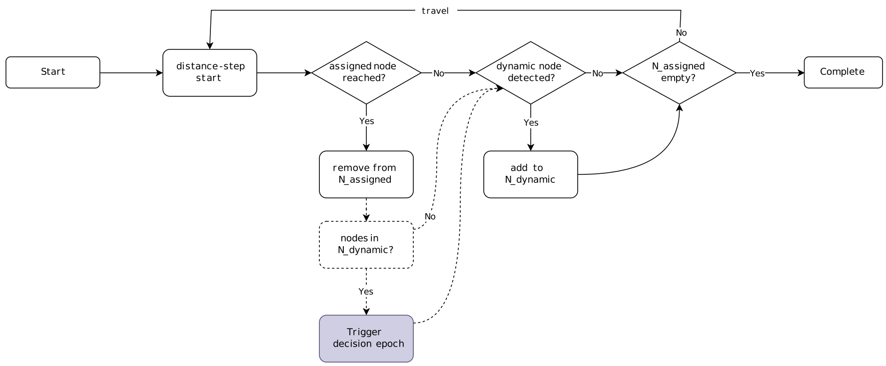

Fig. 2 Simulation process with endogenous trigger. A decision epoch is triggered if an assigned node is reached, AND there are nodes in N_{dynamic}. If no dynamic nodes have been found since the last epoch, no decision epoch occurs. Each distance step checks if an assigned node is reached first, reoptimises if a decision epoch triggered, then checks to detect dynamic nodes. If multiple dynamic nodes have been detected in one distance-step, all will be added to N_{dynamic} for the next distance-step. The decision epoch empties N_{dynamic} and reoptimises routes to include dynamic nodes. (Parts unique to the endogenous trigger are denoted in dotted lines)

3.2 Decisional process: triggers

The simulation checks for dynamic nodes every distance step. The trigger determines when a decision epoch is called. Two different triggers were simulated: exogenous and endogenous. The exogenous case triggers a decision epoch immediately if a dynamic node is detected that distance step (Fig. 1). For endogenous triggers, a decision epoch occurs only after an assigned node has been visited that distance-step (Fig. 2). Because endogenous reoptimisation is triggered based on vehicle status rather than external events, multiple dynamic nodes may be detected while travelling to the assigned node. Hence the number of dynamic nodes in N_{dynamic} is likely to increase over several distance-steps as part of the informational process before being dealt with at the decision epoch.

3.3 Reoptimisation

Once a decision epoch is triggered, the reoptimisation process is the same regardless of how it was triggered. At any given decision epoch k, the system will have the state $S_k = (d_k, (x, y)_k, R_k, N_{\text{assigned}_k}, N_{\text{dynamic}_k})$, where d_k is the distance step associated with decision epoch k, $(x, y)_k$ is the vehicle's coordinates, R_k is the set of current route plans for N_{assigned_k}, and N_{dynamic_k} contains all dynamic nodes detected in the distance between d_k and d_{k-1}. The distance between epochs is determined by when an epoch is triggered.

Once the epoch state is retrieved, reoptimisation occurs; route distance is minimised as the objective function. Using R_k as a basis, new routes are created to assign both N_{dynamic_k} and N_{assigned_k}. Note that reoptimisation is unrestricted: N_{dynamic_k} is not simply inserted into the current routes. Rather, the current routes are used as an initial solution for the optimisation heuristic. Therefore the previously assigned nodes may be re-assigned to different routes based on what is optimal for the problem at epoch k (see Thomas 2011). As the nodes in N_{dynamic_k} are now in the new N_{assigned} after reoptimisation, N_{dynamic} will be empty again at the end of the decisional process.

When the vehicle arrives at an assigned node location, the node is moved from N_{assigned} to N_{visited}. The vehicle travels until all routes are completed and thus all assigned nodes have been visited: $N_{\text{assigned}} = \emptyset$. Because the detection of dynamic nodes is dependent on the vehicle route, the number of total nodes visited—and thus the final value of D—is unknown until the vehicle has completed all routes. A consequence of local information availability is that V_{dynamic} is not guaranteed to be empty, as some dynamic nodes may not have been detected by the end of the solution when the vehicle reaches final distance D.

4 Computational experiments

4.1 Specifications

The simulation and dynamic routing framework was designed modularly, so that any static VRP solver can be used for route planning. For this study, an open-source Guided Local Search (GLS) metaheuristic from the or-tools constraint solver (ver. 7.1) was used. Because computational time was not a consideration for this study, the static solver allocated up to 300 s of computational time for the initial a priori solution. Reoptimisation was allocated up to 2 s per dynamic node. Such generous time allowances were given in order to minimise variance created by the heuristic algorithms used by the static solver. All simulations were run on a 2.3 GHz Intel Xeon Gold 6148/6248 CPU.

4.2 Test instances

Test instances were sourced from the extended Set X CVRP benchmark instances by Uchoa et al. (2017). The Set X benchmarks were chosen over Solomon's historically preferred DVRPTW benchmarks because the geographical distribution instances found in Set X are more extensive and representative of real-life applications. Of the 600 instances in the extended benchmark, the subset of 60 instances designed for geographical distribution analysis were used. The 60 instances are split into three geographical distribution sets of 20 instances: clustered (C), random-clustered (RC) and random (R) distributions, respectively. All benchmark instances come with total distances from their best known solutions. Figure 3 shows the best known solution distances (BKS) for each test instance in the geographical category. The clustered distribution instances tend to have smaller total distances.

All instances have an area of $A = 1000 \text{ m} \times 1000 \text{ m}$, number of nodes $n = 200$, and a random depot location. In order to preserve the node distribution of each geographical category, dynamic nodes are randomly sampled from the existing 200 static nodes in

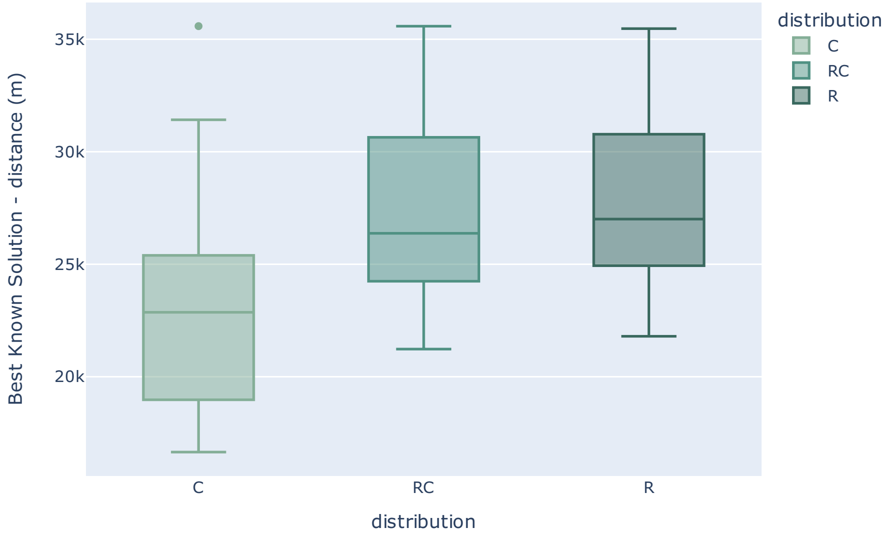

Fig. 3 Best known solutions for the static Set X benchmark instances. Geographical categories: clustered (C), random-clustered (RC), and random (R) each contain 20 unique instances, each with their respective BKS. The range of BKS distances show the variation within each geographical category

Table 1 Simulation configurations

Trigger	Distribution	dod	Detectability
Endogenous	C (20)	0.25	1
Exogenous	RC (20)	0.50	5
	R (20)	0.75	10
			15
			20
			30
			40
			50
			100
			200

Each geographical distribution category has 20 unique instances. Detectability is the radius (meters) of vehicle sight

each benchmark instance. For further information on how node distributions in Set X were generated, please see the technical report by Uchoa et al. (2017).

4.3 Experiment setup

Table 1 shows the levels in each of four attributes tested in the experiments: Exogenous and endogenous trigger; 20 unique instances in each of the C, RC, and

R node distributions; 0.25, 0.5, and 0.75 dod; and 10 levels of detectability. The degrees of dynamism were 0.25, 0.5, and 0.75 to represent low, medium and high dynamism, respectively.

There were 180 attribute combinations (when including detectability levels), applied to each of the 20 unique instances in each geographical category—giving a total of 3600 unique configurations. For each configuration, 120 simulations were run, sampling a different selection of dynamic nodes each run. The global case, which does not have the detectability attribute, had a total of 43,200 simulations; the local information availability case had a total of 432,000 simulations.

Before introducing proximity-dependent nodes, experiments on the classic DVRP were run to create a baseline comparison. Classic DVRPs have global information availability: dynamic nodes will randomly appear, independent of vehicle movement, over the course of the vehicle's routes. The frequency of dynamic nodes appearing per distance step λ followed a Poisson distribution:

$$\lambda = \frac{n \times \text{dod}}{D}$$

where n is the number of nodes in an instance ($n = 200$ for all test instances); dod is the percentage of dynamic nodes, and D is the total distance traveled the dynamic node appearances will be spread over. Because the final distance D of the dynamic instance cannot be known a priori, the benchmark's best known solution distance BKS was used as an estimate for the final distance. The static BKS distance will never be more than the dynamic D, so using the BKS will ensure that most of the dynamic nodes will appear by the time all $V_{\text{a priori}}$ has been visited. Therefore we set $\lambda = \frac{200 \times \text{dod}}{\text{BKS}}$.

For the global information availability case, all dynamic nodes are guaranteed to appear; thus the solution performance was measured by total distance. For local information availability, where the appearance of dynamic nodes are dependent on vehicle proximity, the performance measure was both the proportion of dynamic nodes detected as well as the total distance.

4.4 Performance comparison

To identify the statistically significant attributes that affect the proportion of dynamic nodes detected and the total distance travelled, a Generalised Linear Model (GLM) was used. Let n_{detected} be the number of dynamic nodes detected when under local information availability. Then, $n_{\text{detected}} \sim \text{Binomial}(V_{\text{dynamic}}, P_j)$ where P_j is the proportion of dynamic nodes detected for simulation configuration j, and:

$$\text{logit}(P_j) = \mathbf{X}_\mathbf{j}^\top \beta$$

where $\mathbf{X}$ is the attribute matrix and β is the vector of coefficients. A gamma-distributed model was used to model total distance travelled for global information availability. Starting with a main effect model, interaction effects were then systematically added. Competing models were evaluated by the Akaike Information Criterion

Fig. 4 Total distance travelled vs nodes visited. Visited nodes include both known and detected dynamic nodes

(AIC); given the same fit, the AIC selects the model that minimises complexity and maximises quality of information gained (Akaike 1974).

5 Results

5.1 Exogenous vs. endogenous triggers

Overall, across all simulations and attribute configurations, using an endogenous trigger was statistically superior to an exogenous trigger. In terms of detecting dynamic nodes, there was not much difference between the trigger performances when detectibility was low, or when dod was low and dynamic nodes were clustered. However, as the detectability increased, the proportion of dynamic nodes detected increased at a greater rate for the endogenous trigger compared to the exogenous trigger.

The endogenous trigger was also better at minimising distance travelled. The triggers had similar total distances when the number of nodes visited was small (≤ 50), but past 60+ nodes visited, using the endogenous trigger was statistically advantageous in comparison to the exogenous trigger (Fig. 4).

Although visually Fig. 4 shows the exogenous trigger had a higher range of longer travel distances, statistically the exogenous trigger was actually better at minimising distance when used on the R and RC distribution—albeit only at low dod. At high dod, the exogenous trigger has an overall lengthening affect on the distance travelled (see Table 4 in Appendix).

For the classic DVRP (with global information availability), the endogenous trigger similarly outperformed the exogenous trigger in minimising the total distance

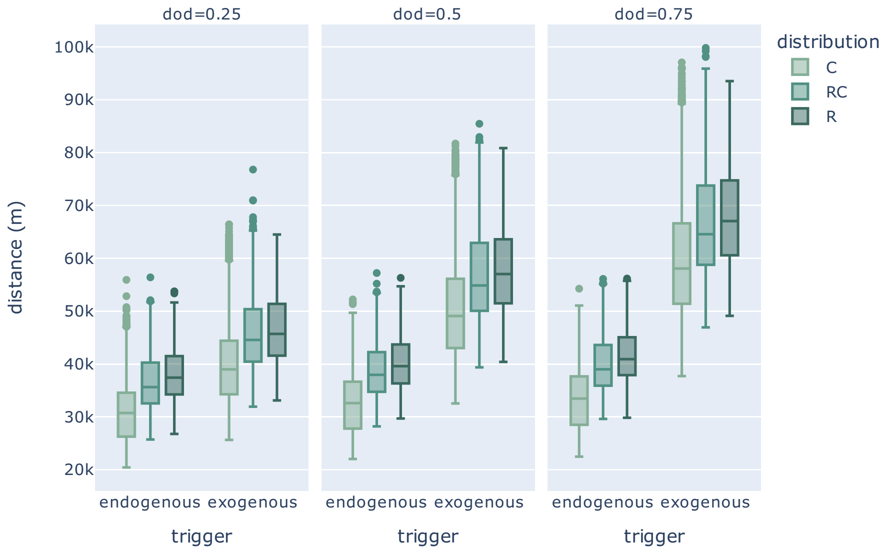

Fig. 5 Comparison of total distance travelled under global information availability. The differences between distances on the C, RC, and R distributions reflect the static BKS baselines shown in Fig. 3

travelled, especially at high dods (Fig. 5). Similar to the local information availability case, the global availability case also saw the exogenous trigger having a *lessening* effect on the total distance (Table 2 in Appendix) when performing on random distributions—even at high dod. This interaction effect with the exogenous trigger was statistically significant, but not necessarily strong enough comparatively against the advantage of using an endogenous trigger overall.

5.2 Effects of dod, node distribution and detectibility

Detectability, dod, geographical distribution, and trigger all had a statistically significant ($p < 0.05$) effect on the proportion of dynamic nodes detected (Table 3 in Appendix). Increasing detectability directly increases the proportion of dynamic nodes detected, with the effect plateauing once the detectability radius exceeds 50 m (Fig. 6).

The degree of dynamism is inversely correlated with the detection rate, with a higher proportion of dynamic nodes detected at lower dod (0.25) compared to high dod (0.75), given fixed trigger and distribution. An implication is that as dod increases, a higher detectability is required to achieve the same number of nodes visited.

Finally, the geographic distribution of nodes had a significant effect on the proportion of dynamic nodes detected. At fixed trigger and detectability, C distributions statistically had the highest dynamic node detection rate, followed by RC, then R (row 6, 7 of Table 3 in Appendix). However, as detectability increased, the detection rate for RC and R distributions became more similar in comparison to C distributions. Thus a higher detectability will be required to achieve the same performance as the distribution goes from clustered to random-uniform, because dynamic nodes will be more spread out.

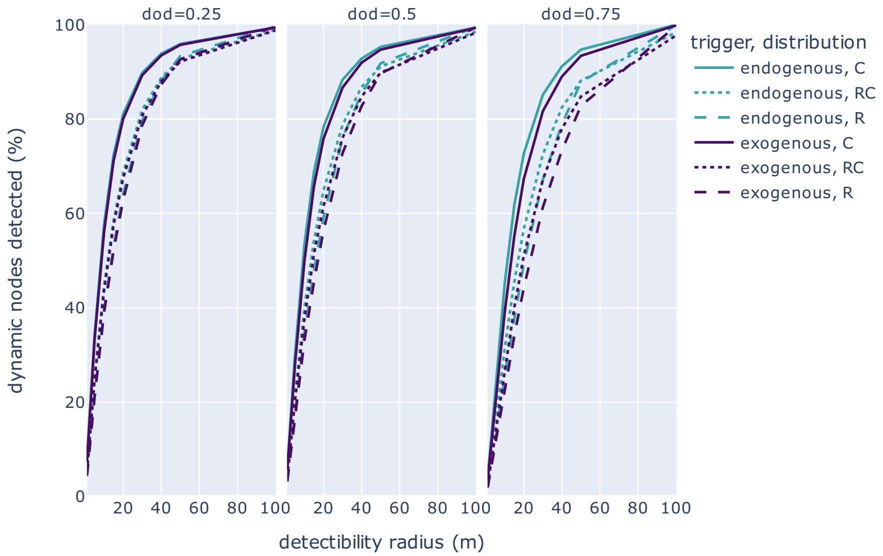

Fig. 6 Average proportion of dynamic nodes detected vs. detectability. Values are averaged over 120 simulations

6 Discussion

6.1 Trigger comparison

The results showed that the endogenous trigger performed better than the exogenous trigger in both detecting dynamic nodes and minimising distance travelled. The exogenous trigger performed best at a low degree of dynamism (dod), which aligns with the findings in Regan et al. (1998) and Ulmer et al. (2017). Ulmer et al. (2017) reasoned an exogenous trigger performs badly at high dod due to reoptimising in quick succession, creating a series of premature, inefficient moves. At low dod, dynamic nodes appear less frequently, so the reoptimised solution is more likely to perpetuate. By this reasoning, an exogenous trigger should perform significantly better on the random distributions as well. The results in both the classic DVRP as well as the DVRP with proximity-dependent nodes confirmed this. Although endogenous trigger was superior overall, perhaps creating a hybrid trigger that incorporates the strengths of both endogenous and exogenous strategies could be a viable approach for problems with a range of geographical node distributions.

The superior performance of the endogenous trigger over exogenous is in line with the findings in Moretti Branchini et al. (2009), but contrary to Ulmer et al. (2017), Lorini et al. (2011); and Ichoua et al. (2000). It is possible that the contradiction is due to a difference in objective function. Most literature which favoured the exogenous trigger were for DVRPs with Time Windows (DVRPTW), whose time-constrained objective functions work better with immediate reoptimastion.

In addition, DVRPTW experiments typically use Solomon's benchmarks or similar, with instances having only one distribution per geographical category (as

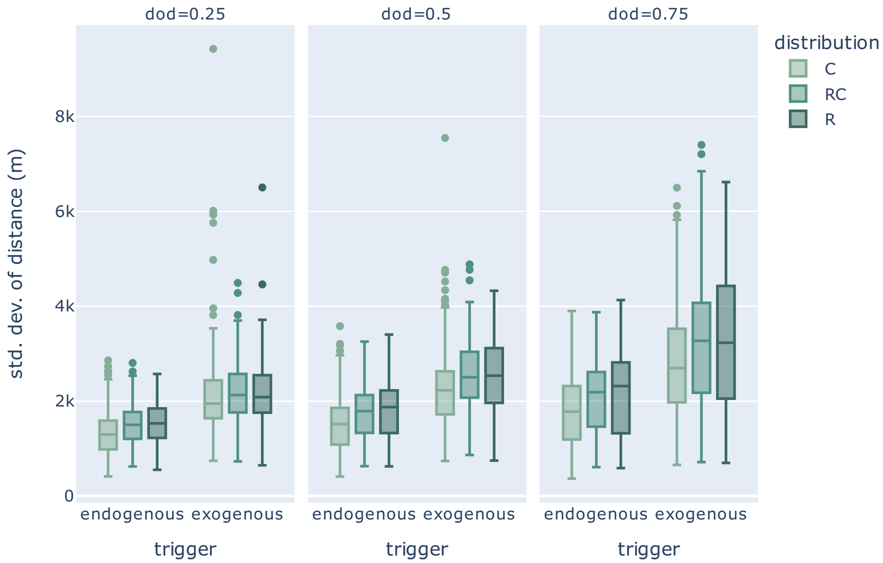

Fig. 7 Standard deviation in distance due to dynamic sampling variation, under local information availability

opposed to Set X's 20 different instances per geographical category). The variation of distribution within a geographical category needs to be considered; the Set X static solutions (Fig. 3) and the results of this study (Figs. 4, 5) demonstrated there can be large variability within a geographical category, and a clustered instance can easily have the same results as a random instance. Experiments that only use one instance per geographical category have a high probability of bias, which impacts the applicability of the strategic decisions recommended.

6.2 Variability

To prevent bias in high variability studies, it is important that results are statistically proven. Note that there was a large overlap between the endogenous and exogenous solutions in Fig. 4 due to high variance. Figure 7 shows the standard deviation in the distance over the 120 simulations for any given experimental configuration. As the dod increased, trigger performance became more varied, especially for the exogenous trigger. This is because a high dod increases the sampling pool from which dynamic nodes are chosen, which creates more variance in the initial distribution of known nodes. The wide range of standard deviations suggests that the exogenous trigger was more sensitive to the distribution of nodes, as opposed to the more robust endogenous trigger.

There were two main sources of variability in this study: the geographical variability of node placement in the 20 instances of each geographical category (Fig. 3); and the geographical variability in dynamic nodes sampled for each of the 120 simulations (Fig. 7). There was also an element of variability from using a heuristic search to reoptimise routes, but it was considered negligible. A potential way to reduce variability in future is to use a finer categorisation for geographical distribution, instead of broadly categorising a wide range of instances into only three distribution sets. While this study has statistically significant results for Set X, more research is required on a wide variety of problems to get a more robust understanding of the conditions which affect trigger suitability. An interesting avenue of future work in this area is the degree of structural diversity introduced by Ferrucci et al. (2013), which is a continuous measure of the geographical variability of dynamic nodes.

6.3 Environmental variables

There are several implications as to how the environmental variables affected the total distance travelled and detection rate. When comparing the global and local information availability cases that visited all 200 nodes, distance travelled in local case tended to reach lower values than for the global case—no doubt due to the differing mechanisms behind dynamic nodes appearance. Detecting 200 nodes in the local case would mean a high detectability, more dynamic nodes detected earlier on, and thus slightly more efficient routing. On the other hand, the global case uniformly spreads out dynamic node appearance throughout the travel period as per a Poisson process, so routing would have been slightly less efficient.

The detection rate in the local case was logarithmically dependent on the detectibility (Fig. 6). In practical terms, this means that there is an optimum detectibility-to-detection ratio.

For example, a practitioner prototyping a new autonomous drone for agricultural use may need to choose a vision sensor (e.g., LiDAR, radar, sonar, or camera), each with a different range of vision and associated cost. Given a similar geographical distribution to Uchoa's clustered benchmarks, and a dod of 25%, the practitioner will know from our results in Fig. 6 that the ideal detectability range of 30 ms will detect 90% of the dynamic nodes on average. However, if their application has a dod of 75%, then they should increase their detectibility approximately 40 ms for the same results. Of course, there is variability, but by inputting the known environmental conditions, the practitioner can obtain a starting estimate on the range of sight most suitable for their application context.

6.4 Future works

This study only investigated the simplest case of DVRP with a single vehicle and two trigger types; there are many extensions to the DVRP with proximity dependent nodes. More realistic applications may include online variants (for when no a-priori information is available), or parallelised solutions for a fleet of decentralised autonomous agents. To be applicable across problems of different sizes, it may be useful in future to define the detectability as a unitless parameter, such as a percentage of the total problem area. Furthermore, hybrid triggers that combine the advantages of quick, local insertions and more longer-term global reoptimisation may be investigated.

For applications that must guarantee all dynamic nodes are detected, the detectability needs to be increased to a high enough level. For applications with an unchangable detectibility, a possible solution would be to send the vehicle outside its shortest routes to search unexplored areas not visible from its route path. This exploration approach would be like a dynamic stochastic variation of the Aerial Surveillance Problem (Karasakal 2016) or the Maximum Cover Shortest Path Problem (Boffey et al. 1995).

Future work will involve tackling the strong coupling between the decisional process and the problem state, specifically by integrating anticipatory methods into the solution approach. An anticipatory approach will require an objective function that prioritises information collection in addition to optimising route distances.

7 Conclusion

In this study, the concept of proximity-dependent nodes was introduced, and detectability was defined as an attribute of the Dynamic Vehicle Routing Problem with proximity-dependent nodes. Endogenous and exogenous reoptimisation strategies were compared for the classic DVRP as well as the DVRP with proximity-dependent nodes. Although the problem variant is unusual, the reoptimisation strategies were tested on the well-known Uchoa benchmarks to enable easy comparison and encourage further development in this area.

The results of the computational experiments showed that the endogenous trigger was better in minimising distance and detecting dynamic nodes. Results indicate that an exogenous trigger should only be considered when the dynamism is low, or the detectability is very low. Extra care should be taken when measuring performance of an exogenous trigger on problems with high dynamism, as solution quality under these conditions can be highly variable.

Furthermore, the effects of the degree of dynamism and geographical node distribution on the detection rate was investigated. Results indicated that the detection rate was lower when the dynamism of a problem is high, and the geographical distribution of nodes is not clustered. Thus, a higher detectibility radius is required to reach the same level of detection as dynamism increases and nodes are more spread out. Overall, the proportion of dynamic nodes detected increased logarithmically with detectability, indicating that a practitioner can in fact tune for the optimal level of detectability, based on the problem conditions.

Appendix A: Generalised linear models

See Tables 2, 3 and 4.

Table 2 Generalised Linear Model predicting total distance travelled with global information availability

Attribute	Coefficient	Std. err	t value	$Pr(>\|t\|)$	Significance
Intercept	10.356366	0.003273	3164.064	< 2e−16	***
Exogenous	0.248031	0.004629	53.583	< 2e−16	***
R	0.187867	0.004629	40.586	< 2e−16	***
RC	0.147536	0.004629	31.873	< 2e−16	***
dod = 0.75	0.078139	0.004629	16.881	< 2e−16	***
dod = 0.5	0.049691	0.004629	10.735	< 2e−16	***
dod = 0.75:exogenous	0.319622	0.006546	48.825	< 2e−16	***
dod = 0.5:exogenous	0.178824	0.006546	27.317	< 2e−16	***
dod = 0.75:R	0.011758	0.006546	1.796	0.072471	
dod = 0.75:RC	0.011456	0.006546	1.750	0.080116	
dod = 0.5:R	0.005793	0.006546	0.885	0.376167	
dod = 0.5:RC	0.008341	0.006546	1.274	0.202619	
Exogenous:R	− 0.042529	0.006546	− 6.497	8.3e− 11	***
Exogenous:RC	− 0.024998	0.006546	− 3.819	0.000134	***
dod = 0.75:exogenous:R	− 0.033005	0.009258	− 3.565	0.000364	***
dod = 0.75:exogenous:RC	− 0.034648	0.009258	− 3.743	0.000182	***
dod = 0.5:exogenous:R	− 0.019120	0.009258	− 2.065	0.038898	*
dod = 0.5:exogenous:RC	− 0.021680	0.009258	− 2.342	0.019198	*

Distribution = C, dod = 0.25, trigger = endogenous are set to 0 as the baselines coefficients for this GLM. Variables with a p value of 0.05 or less (* to ***) are considered statistically significant. Significance codes: ***:0.001, **:0.01, *:0.05, .:0.1,‘ ’:1 p value

Table 3 Generalised Linear Model predicting proportion of detected dynamic nodes

Attribute	Coefficient	Std. err	t value	$Pr(>\|t\|)$	Significance
(Intercept)	− 1.2555922	0.0176533	− 71.125	< 2e−16	***
Detectability	0.1231071	0.0009809	125.506	< 2e−16	***
Exogenous	− 0.0314921	0.0143730	− 2.191	0.0284	*
dod = 0.5	− 0.1481174	0.0174871	− 8.470	< 2e−16	***
dod = 0.75	− 0.4384530	0.0177651	− 24.681	< 2e−16	***
RC	− 0.2371020	0.0176519	− 13.432	< 2e−16	***
R	− 0.4454384	0.0179422	− 24.826	< 2e−16	***
Detectability:exogenous	− 0.0057908	0.0006905	− 8.387	< 2e−16	***
Detectability:dod = 0.5	− 0.0047948	0.0008921	− 5.375	7.67e−08	***
Detectability:dod = 0.75	− 0.0110035	0.0008568	− 12.843	< 2e−16	***
Detectability:RC	− 0.0215280	0.0009229	− 23.327	< 2e−16	***
Detectability:R	− 0.0209751	0.0009168	− 22.878	< 2e−16	***

Distribution = C, dod = 0.25, trigger = endogenous are set to 0 as the baselines coefficients for this GLM. Variables with a p value of 0.05 or less are considered statistically significant. Significance codes: ***:0.001 **:0.01, *:0.05, .:0.1, ‘ ’:1 p value

Table 4 Generalised Linear Model predicting distance travelled with local information availability

Attribute	Coefficient	Std. err	t value	$Pr(>\|t\|)$	Significance
(Intercept)	10.203817	0.001356	7522.648	< 2e−16	***
Exogenous	0.143924	0.001918	75.028	< 2e−16	***
R	0.173715	0.001918	90.559	< 2e−16	***
RC	0.136451	0.001918	71.133	< 2e−16	***
dod = 0.5	0.035233	0.002059	17.115	< 2e−16	***
dod = 0.75	0.021201	0.002199	9.639	< 2e−16	***
Exogenous:R	− 0.029415	0.002713	− 10.843	< 2e−16	***
Exogenous:RC	− 0.020525	0.002713	− 7.566	3.87e−14	***
Exogenous:dod = 0.5	0.093167	0.002921	31.896	< 2e−16	***
Exogenous:dod = 0.75	0.202052	0.003151	64.131	< 2e−16	***
R:dod = 0.5	0.038808	0.002999	12.941	< 2e−16	***
R:dod = 0.75	0.041727	0.003275	12.741	< 2e−16	***
RC:dod = 0.5	0.019525	0.002963	6.590	4.39e−11	***
RC:dod = 0.75	0.018321	0.003226	5.679	1.35e−08	***
Exogenous:R:dod = 0.5	0.021940	0.004259	5.151	2.59e−07	***
Exogenous:R:dod = 0.75	0.032398	0.004705	6.885	5.77e−12	***
Exogenous:RC:dod = 0.5	0.015543	0.004207	3.695	0.00022	***
Exogenous:RC:dod = 0.75	0.018932	0.004624	4.094	4.23e−05	***

To compare across all dods equally, the model is based on distance travelled when 150–200 nodes are visited. Distribution = C, dod = 0.25, trigger = endogenous are set to 0 as the baselines coefficients for this GLM. Variables with a p value of 0.05 or less are considered statistically significant. Significance codes: ***:0.001, **:0.01, *:0.05, .:0.1, ‘ ’:1 p value

Acknowledgements The authors are grateful to CVRPLIB community for promptly providing the extended benchmark files and answering questions. Furthermore, thanks are due the referees for their valuable comments. This work was supported by the Australian Government Research Training Program Scholarship.

Funding Open Access funding enabled and organized by CAUL and its Member Institutions.

Data availability Raw experimental data and other additional data used in this paper is publically available on Github via: https://zenodo.org/badge/latestdoi/638083008. Set X extended CVRP benchmarks are publically available at CVRPLIB: http://vrp.galgos.inf.pucrio.br/index.php/en/.

References

Akaike H (1974) A new look at the statistical model identification. IEEE Trans Autom Control 19(6):716–723. https://doi.org/10.1109/TAC.1974.1100705

Al-Kanj L, Powell WB, Bouzaiene-Ayari B (2016) The information-collecting vehicle routing problem: stochastic optimization for emergency storm response. arXiv:1605.05711

Boffey B, García FRF, Laporte G, Mesa JA, Pelegrín BP (1995) Multiobjective routing problems. Top 3(2):167–220. https://doi.org/10.1007/BF02568585

Boysen N, Schwerdfeger S, Weidinger F (2018) Scheduling last-mile deliveries with truck-based autonomous robots. Eur J Oper Res 271(3):1085–1099. https://doi.org/10.1016/j.ejor.2018.05.058

Braekers K, Ramaekers K, Van Nieuwenhuyse I (2016) The vehicle routing problem: state of the art classification and review. Comput Ind Eng 99:300–313. https://doi.org/10.1016/j.cie.2015.12.007

Bsaybes S, Quilliot A, Wagler AK (2019) Fleet management for autonomous vehicles using flows in time-expanded networks. Top 27(2):288–311. https://doi.org/10.1007/s11750-019-00506-4

Chen C, Demir E, Huang Y (2021) An adaptive large neighborhood search heuristic for the vehicle routing problem with time windows and delivery robots. Eur J Oper Res 294(3):1164–1180. https://doi.org/10.1016/j.ejor.2021.02.027

Dror M, Laporte G, Trudeau P (1989) Vehicle routing with stochastic demands: properties and solution frameworks. Transp Sci 23(3):166–176. https://doi.org/10.1287/trsc.23.3.166

Eglese R, Zambirinis S (2018) Disruption management in vehicle routing and scheduling for road freight transport: a review. Top 26(1):1–17. https://doi.org/10.1007/s11750-018-0469-4

Ferrucci F (2013) Pro-active dynamic vehicle routing: real-time control and request-forecasting approaches to improve customer service, 1st edn. Contributions to Management Science. Physica-Verlag Heidelberg, Heidelberg. https://doi.org/10.1007/978-3-642-33472-6

Ferrucci F, Bock S (2015) A general approach for controlling vehicle en-route diversions in dynamic vehicle routing problems. Transp Res Part B Methodol 77:76–87. https://doi.org/10.1016/j.trb.2015.03.003

Ferrucci F, Bock S, Gendreau M (2013) A pro-active real-time control approach for dynamic vehicle routing problems dealing with the delivery of urgent goods. Eur J Oper Res 225(1):130–141. https://doi.org/10.1016/j.ejor.2012.09.016

Gendreau M, Guertin F, Potvin JY, Taillard É (1999) Parallel Tabu search for real-time vehicle routing and dispatching. Transp Sci 33(4):381–390. https://doi.org/10.1287/trsc.33.4.381

Gendreau M, Jabali O, Rei W (2016) Future research directions in stochastic vehicle routing. Transp Sci 50(4):1163–1173

Hyland M, Mahmassani HS (2017) Dynamic autonomous vehicle fleet operations: optimization-based strategies to assign AVs to immediate traveler demand requests. Transp Res Part C Emerg Technol 92:278–297. https://doi.org/10.1016/j.trc.2018.05.003

Ichoua S, Gendreau M, Potvin JY (2000) Diversion issues in real-time vehicle dispatching. Transp Sci 34(4). https://doi.org/10.1287/trsc.34.4.426.12325

Karasakal O (2016) Minisum and maximin aerial surveillance over disjoint rectangles. Top 24(3):705–724. https://doi.org/10.1007/s11750-016-0416-1

Larsen A (2000) The dynamic vehicle routing problem, Phd. Technical University of Denmark

Larsen A, Madsen OBG, Solomon MM (2007) Classification of dynamic vehicle routing systems. Springer, Boston, pp 19–40. https://doi.org/10.1007/978-0-387-71722-7_2

Lorini S, Potvin JY, Zufferey N (2011) Online vehicle routing and scheduling with dynamic travel times. Comput Oper Res 38(7):1086–1090. https://doi.org/10.1016/j.cor.2010.10.019

Lund K, Madsen OBG, Rygaard JM (1996) VRPs with varying degrees of dynamism. Tech. rep, Institute of Mathematical Modelling, Lyngby

Margolis JT, Song Y, Mason SJ (2022) A Markov decision process model on dynamic routing for target surveillance. Comput Oper Res 141:105699. https://doi.org/10.1016/j.cor.2022.105699

Moretti Branchini R, Amaral Armentano V, Løkketangen A (2009) Adaptive granular local search heuristic for a dynamic vehicle routing problem. Comput Oper Res 36(11):2955–2968. https://doi.org/10.1016/j.cor.2009.01.014

Pillac V, Gendreau M, Guéret C, Medaglia AL (2013) A review of dynamic vehicle routing problems. Eur J Oper Res 225(1):1–11. https://doi.org/10.1016/j.ejor.2012.08.015

Poikonen S, Wang X, Golden BL (2017) The vehicle routing problem with drones: extended models and connections. Networks 70(1):34–43. https://doi.org/10.1002/net.21746
Psaraftis HN (1995) Dynamic vehicle routing: status and prospects. Ann Oper Res 61:143–164
Psaraftis HN, Wen M, Kontovas CA (2016) Dynamic vehicle routing problems: three decades and counting. Networks 67:3–31. https://doi.org/10.1002/net.21628
Regan AC, Mahmassani HS, Jaillet P (1998) Evaluation of dynamic fleet management systems simulation framework. Transp Res Rec 1645(1645):176–184. https://doi.org/10.3141/1645-22
Ritzinger U, Puchinger J, Hartl RF (2016) A survey on dynamic and stochastic vehicle routing problems. Int J Prod Res 54(1):215–231. https://doi.org/10.1080/00207543.2015.1043403
Savelsbergh MWP, Van Woensel T (2016) 50th anniversary invited article-city logistics: challenges and opportunities. Transp Sci 50(2):579–590. https://doi.org/10.1287/trsc.2016.0675
Solomon MM (1987) Algorithms for the vehicle routing and scheduling problems with time window constraints. Oper Res 35(2):254–265. https://doi.org/10.1287/opre.35.2.254
Thomas BW (2011) Dynamic vehicle routing. In: Wiley Encyclopedia of Operations Research and Management Science. Wiley, Hoboken. https://doi.org/10.1002/9780470400531.eorms0278
Uchoa E, Pecin D, Pessoa A, Poggi M, Vidal T, Subramanian A (2017) New benchmark instances for the capacitated vehicle routing problem. Eur J Oper Res 257(3):845–858. https://doi.org/10.1016/j.ejor.2016.08.012. arXiv:1405.7020v1
Ulmer MW, Heilig L, Voß S (2017) On the value and challenge of real-time information in dynamic dispatching of service vehicles. Bus Inf Syst Eng 59(3):161–171. https://doi.org/10.1007/s12599-017-0468-2
Van Woensel T, Kerbache L, Peremans H, Vandaele N (2008) Vehicle routing with dynamic travel times: a queueing approach. Eur J Oper Res 186(3):990–1007. https://doi.org/10.1016/j.ejor.2007.03.012

Authors and Affiliations

Tiria Andersen[1] · Shaun Belward[1] · Mangalam Sankupellay[1] · Trina Myers[2] · Carla Chen[1]

Shaun Belward
shaun.belward@jcu.edu.au

Mangalam Sankupellay
mangalam.sankupellay@gmail.com

Trina Myers
trina.myers@qut.edu.au

[1] College of Science and Engineering, James Cook University, 1 James Cook Dr, Townsville, QLD 4811, Australia

[2] Faculty of Science, Queensland University of Technology, 2 George Street, Brisbane, QLD 4000, Australia