Supporting Humans in Solving Multi-UAV Dynamic Vehicle Routing Problems

Klein Koerkamp, N.W.; Borst, Clark; Mulder, Max; van Paassen, Rene

DOI
10.1016/j.ifacol.2019.12.088

Publication date
2019

Document Version
Final published version

Published in
IFAC-PapersOnLine

Available online at www.sciencedirect.com

ScienceDirect

IFAC Papers Online CONFERENCE PAPER ARCHIVE

IFAC PapersOnLine 52-19 (2019) 359–364

Supporting Humans in Solving Multi-UAV Dynamic Vehicle Routing Problems

N. W. Klein Koerkamp C. Borst [1] **Max Mulder**
M. M. van Paassen

Control and Simulation, Faculty of Aerospace Engineering, TU Delft, 2629 HS, Delft, The Netherlands

Abstract: Real-time optimization of Vehicle Routing Problems (VRP) during mission operations raises concerns regarding obtaining a solution within a reasonable timeframe, especially in domains where operations cannot easily be paused and the number of control parameters is high. Humans, however, are *heuristic* problem solvers and could potentially complement VRP algorithms in providing quickly a workable and safe solution from which the algorithms can further find the optimum. In this study, a visual interface was developed and evaluated aiming to support humans in manually solving a dynamic VRP in which they needed to solve various simulated payload delivery missions, featuring multiple Unmanned Aerial Vehicles, under failure conditions. Experiment results ($n = 16$) indicate that the interface enabled the majority of participants to quickly solve the perturbed scenarios, although not always in the most efficient way. Interestingly, participants experienced most difficulty in solving the seemingly easier scenarios, featuring less customers and a relatively low number of vehicles compared to the more complex scenarios.

Keywords: Human-machine interface, vehicle routing problem, unmanned aerial vehicles.

1. INTRODUCTION

Vehicle routing problems (VRP) are at the core of many logistics applications. With the current focus on just-in-time logistics and data-driven analysis techniques, cost-efficient routing of a fleet of vehicles plays an important role in many industries (Lin et al., 2014). Determining efficient routes for each vehicle in the fleet, such that the overall mission goal is achieved within distance, capacity and time constraints, defines the VRP (Laporte, 2007). Although much attention is given in literature to obtaining routes for static problems, in many real-life situations these problems become dynamic due to changing customer locations, vehicle failures or uncertain service and travel times caused by perturbations during mission operations (Laporte, 1992; Pillac et al., 2013).

In solving dynamic VRPs, various challenges can be identified. First, the optimization problem needs to be explicitly formulated. Although a wide range of VRP types has been studied in literature, constructing an algorithm that takes into account all possible disturbances and stochastic properties of a real-life application is challenging (Psaraftis et al., 2016). When not taking into account all constraints inflicted by these disturbances and properties, a theoretical optimal solution will not transfer well to real-life operations. Second, due to the complexity of the VRP optimization problem and current computational limitations, finding an optimal solution might take a long time (in the order of hours or sometimes even days) (Toth and Daniele, 2014). Although this might not be a problem for generating schedules well before mission execution, it is a concern when a need for re-optimization arises due to perturbations during mission operations. Third, for some scenarios there might not even be a solution. For example, for over-constrained problems, VRP algorithms might be unable to find a solution at all, without first performing some kind of constraint relaxation (Lau et al., 2003).

A possible way to overcome these challenges is by having humans solving these problems as they are creative problem solvers who can adapt to novel situations. Previous research on human control performance in solving Traveling Salesman Problems (TSP) shows good human performance in creating optimal routes purely based on a visual representation of the customer locations and vehicle routes (Anderson et al., 2000; Vickers et al., 2001; Scott et al., 2002; Dry et al., 2006; Tütüncü et al., 2009; Macgregor and Chu, 2011). By leveraging these human visual pattern recognition qualities, good performance in solving real-time dynamic VRPs might be achieved.

This study comprises both the design and evaluation of a visual interface for solving a dynamic VRP under perturbations. As such, it aims to go a step further than previous research by focusing on human control performance in solving more complex TSP scenarios. A human-in-the-loop experiment focused on a multi-UAV (Unmanned Aerial Vehicle) payload delivery mission with vehicle failures. The task of the operator was to perform real-time perturbation management to ensure mission success. This application has been chosen due to the inherent time pressure introduced with flight operations, where vehicles have limited payload and endurance due to battery capacity.

[1] E-mail: c.borst@tudelft.nl

10.1016/j.ifacol.2019.12.088

2. THE VEHICLE ROUTING PROBLEM

The generic family of VRPs can be defined as follows:

"Given a set of transportation requests and a fleet of vehicles, determine a set of vehicle routes to perform all transportation requests with the given vehicle fleet at minimum cost; in particular, decide which vehicle handles which requests in which sequence so that all vehicle routes can be feasibly executed." (Toth and Daniele, 2014)

Various types of VRPs exist, each with its own specific set of constraints (Toth and Daniele, 2014). Under consideration for this research is the Distance-Constrained Capacitated Vehicle Routing Problem (DCVRP) with resource constraints at the depot, which, in addition to the generic VRP attributes, includes a distance constraint for each vehicle, fuel limitations, a capacity constraint, payload capacity, and a depot capacity limit (i.e., only a finite number of vehicles is allowed to depart and arrive at the depot simultaneously).

More formally, the transportation requests in the DCVRP consist of the distribution of goods from a single depot, denoted as point 0, to customers, which are defined as a set of n other points, with $N = \{1, 2, ..., n\}$. Figure 1 depicts a graphical overview of the generic VRP. The customer demand, $q_i \geq 0$, is defined as the number of goods that needs to be delivered to the customer $i \in N$. The fleet of vehicles, defined as $K = \{1, 2, ..., |K|\}$, that is used to distribute the goods is assumed homogeneous. The homogeneity entails that the $|K|$ vehicles in the fleet have identical distance constraints, capacity $Q > 0$ constraints, and associated cost. Each vehicle starts at the depot, delivers goods to a subset of customers $S \subset N$ visiting customer locations only once, then returns to the depot. When traveling from customer i to customer j, the vehicle incurs the travel cost c_{ij}. Cost is assumed symmetric, where the cost of traveling from i to j is equal to the cost of traveling from j to i.

A dynamic VRP is characterized by perturbations during execution of the original optimized plan, such that re-optimization is necessary (Ercan and Gencer, 2018). Sources of dynamism can include a change in customer demands, increases in travel time or distance due to required re-routes of vehicles and breakdowns of vehicles. In this article, the focus of dynamism is solely put on vehicle failures. In such situations, one or more vehicles can

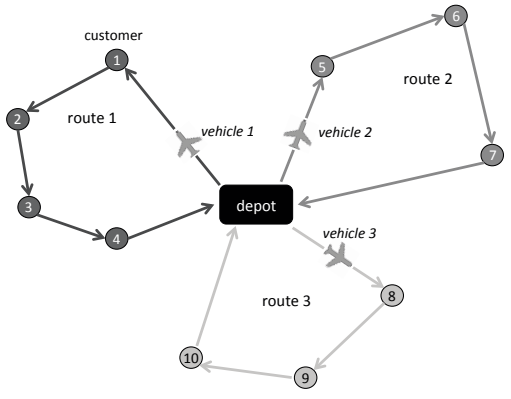

Fig. 1. Graphical representation of the VRP.

breakdown, leaving a number of customers unserved. The remaining vehicles in the fleet should then be used to serve those customers and thus complete the overall mission, requiring each operative vehicle to have a surplus in payload. The consequence is that deviation from the original plan is required and a new optimal, but revised, plan must be found, all within the vehicle fuel limitations (affecting vehicle endurance) and depot capacity constraints.

Although many algorithms, both exact and (meta)heuristic, already exist to optimize and approximate solutions to dynamic VRPs, it is still a topic that is heavily researched (Ercan and Gencer, 2018; Braekers et al., 2016; Psaraftis et al., 2016). Not only does it draw attention because of its notorious difficulty as a combinatorial optimization problem, but also because of its practical relevance (Toth and Daniele, 2014).

3. INTERFACE DESIGN

3.1 Scope and Information Requirements

In this paper the dynamic VRP under consideration corresponds to a DCVRP where the degree of dynamism is defined in terms of vehicle failure. Instead of developing an algorithm that deals with this situation, we focus on developing an interface that supports humans in solving it.

The high-level constraints in the mission are: UAV payload capacity limits, the UAV flight time limit (due to limited battery capacity), and the depot capacity limit. In the current study, communication range (both related to ground station and UAV), airspace restrictions, vehicle separation requirements (with respect to both terrain and other vehicles), weather (such as wind), and UAV flight performance characteristics were left outside the scope.

Given the mathematical formulation of the DCVRP and previous research on human control performance in traveling salesman problems, the following information would be critical to facilitate successful human control performance: 1) a map of geographical customer locations and their demands, 2) vehicle payload capacity, 3) vehicle routes and endurance (based on battery capacity), 4) temporal arrival schedule of the vehicles and 5) depot capacity limit.

3.2 Layout, Structure and Functionality

In designing a visual interface for the dynamic DCVRP control problem considered in this study, inspiration was taken from research in Air Traffic Control that focused on aircraft spatio-temporal arrival management (De Wit et al., 2014) and en-route perturbation management (Klomp et al., 2015).

Figure 2 provides a representation of the layout and structure of the designed direct manipulation interface for a simple scenario. The scenario consists of three UAVs delivering payload to six customers from a single depot location. The interface consists of three separate views, namely the map (A), payload (B), and timeline view (C), where the map view presents information from a spatial perspective and the timeline view presents information from a temporal perspective. The red zone in the timeline view represents the depot capacity constraint.

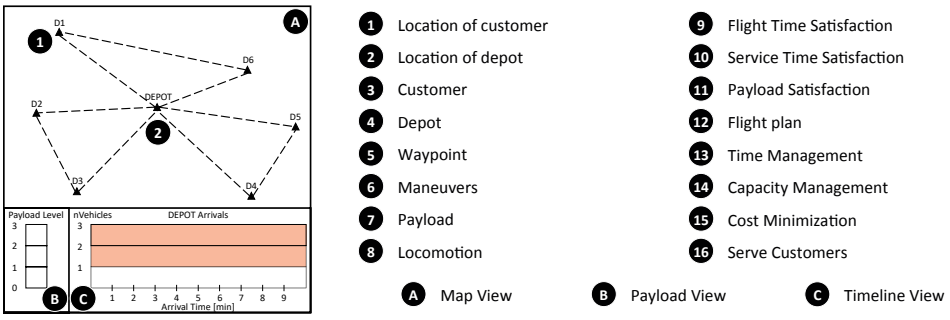

(a) t_0: Mission overview

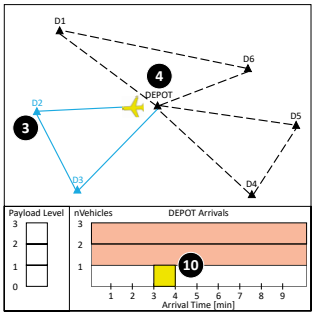

(b) t_1: First UAV launched

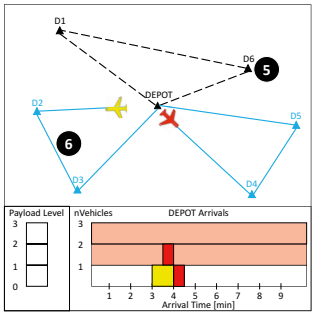

(c) t_2: Second UAV launched

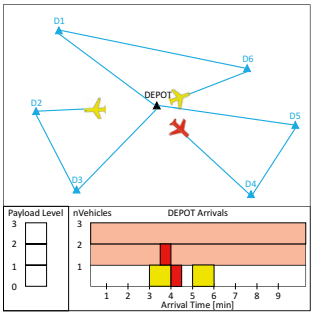

(d) t_3: Third UAV launched

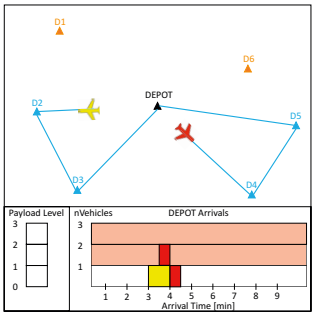

(e) t_4: UAV failure

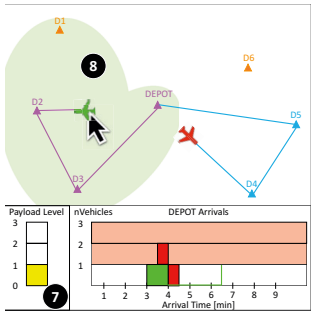

(f) t_5: Select UAV

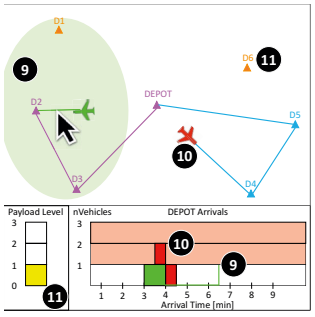

(g) t_6: Select flight plan leg

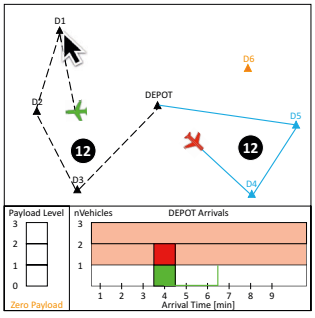

(h) t_7: Add customer D1

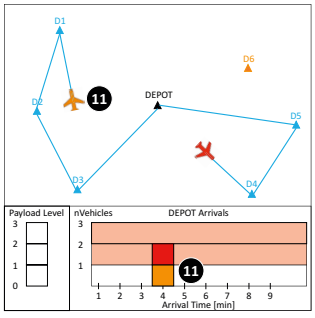

(i) t_8: Confirmed flight plan

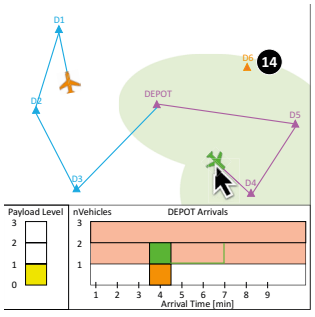

(j) t_9: Select UAV

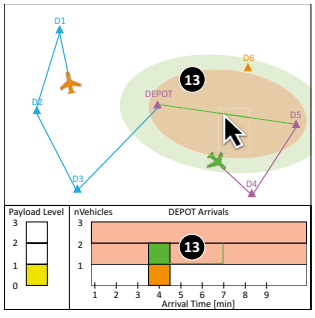

(k) t_{10}: Select flight plan leg

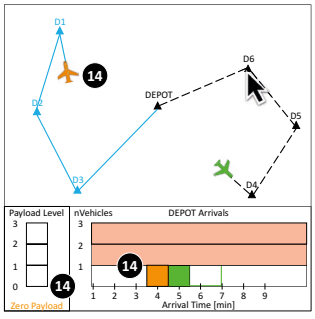

(ℓ)t$_{11}$: Add customer D6

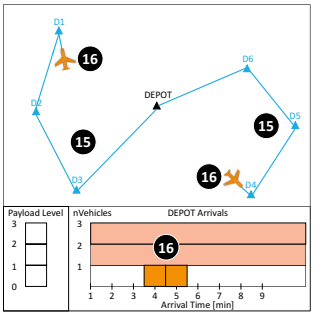

(m) t_{12}: Confirmed flight plan

Fig. 2. Step-by-step overview of the interface for a simple scenario.

At t_0 (Figure 2a) the interface provides a mission overview by displaying customer (1) and depot (2) locations and the initial optimized flight plans (12) for each UAV. The flight plans of vehicles which have not yet left the depot are drawn using a dashed line to differentiate from the vehicles that already have. The first UAV is launched from the depot (4) and flies to its first customer (3) at time t_1 (Figure 2b). The arrival time at the depot and corresponding service time is now also indicated (10) for this vehicle.

The color of the UAV icon and the arrival time block correspond to the payload level of the vehicle, where bright yellow is used when all payload is available, dark yellow when payload capacity is reduced and amber is used when none is available. At t_2 (Figure 2c), the second UAV is launched, of which its arrival time overlaps with the first UAV launched and hence exceeds the depot capacity. Therefore, both the arrival time block in the timeline view as well as the UAV icon in the map view are colored red to assist in activating the operator to identify the problem and take action. Also, the vehicle's flight plan consists of waypoints (5) at the customer locations as well as any intermediate waypoints that the operator my include for path stretching purposes. The future maneuvers a UAV will make are depicted by the guidance reference (6). The third and final UAV is launched at t_3 (Figure 2d).

At t_4 (Figure 2e), one of the vehicles fails, disappearing from the interface and resulting in two customers not being served. One vehicle is selected at time t_5 (Figure 2f), indicated by the green coloring of the UAV icon and the arrival time block. If a vehicle is selected, the payload window indicates the payload available (7). An envelope (8) around the guidance reference indicates what locations can be reached given the vehicle's battery status. At t_6 (Figure 2g), a flight plan leg is selected and the corresponding flight time constraint is indicated (9) both in the map view and in the timeline view, where the vertical line indicates the maximum flight time. The red UAV icon and the arrival time overlap (10) indicate service time issues. The payload level of the vehicles in the fleet and the unvisited customers (11) yield information on the payload satisfaction. Customer D1 is included in the flight plan at t_7 (Figure 2h), where the updated flight plan (12) is indicated with a dashed line and the arrival time is updated. At t_8 (Figure 2i) the modified flight plan is confirmed and the UAV icon and arrival time block (11) changes color to amber in order to indicate no more payload is available after visiting all assigned customers.

At t_9 (Figure 2j), the other UAV has sufficient payload capacity (14) to cover the remaining customer and is selected. Once a flight plan leg is selected at time t_{10} (Figure 2k), in addition to the flight time constraint, the required delay to solve the depot arrival time overlap is visualized (13). This combination gives an integrated overview of time management. At t_{11} (Figure 2ℓ), the modified flight plan is shown. Also, it can be observed that all UAVs in the fleet have now used up their full payload capacity (14). Finally, at t_{12} (Figure 2ℓ) the flight plan is confirmed and the solution to the scenario is visible. Not only are all customers served (16), but this is also achieved using efficient routing of the vehicles (15), while adhering to all applicable constraints.

4. HUMAN-IN-THE-LOOP EXPERIMENT

4.1 Independent Variables

The experiment design consisted of two within-participant independent variables, namely:

(1) Payload capacity (P): The payload capacity of a single UAV, serving as a metric for DCVRP problem size, consisted of four levels: 4, 5, 6 and 7 payload items. All vehicles in a scenario have the same payload capacity.
(2) Perturbation severity (F): The perturbation severity dictates how many UAVs will fail in the scenario, consisting of two levels: single and double failure. All vehicle failures always occurred simultaneously after five seconds into the scenario.

4.2 Scenarios

Participants were asked to mitigate the effects caused by UAV failures during several multi-UAV payload delivery missions under the eight different experiment conditions, see Table 1. The vehicle failures resulted in unassigned customer locations and the task of the participant was to include the unassigned locations into the flight plans of the remaining vehicles, while satisfying all constraints (flight time, payload capacity and depot capacity).

Table 2 lists the number of customer locations and the number of vehicles per condition, both of which are uniquely dictated by the payload capacity, number of vehicle failures, and payload margin. In Figure 3 two example scenarios are shown.

All scenarios lasted six minutes. In every scenario, UAVs were deployed in batches from the depot every thirty seconds (equal to the depot service time), with the batch size equaling the depot capacity. Only lateral control was available, by means of flight plan waypoint modification. Any control actions taken by the participant could influence the solution space later in the scenario.

Table 1. Experiment Conditions

	Payload 4	Payload 5	Payload 6	Payload 7
1 Failure	F1P4	F1P5	F1P6	F1P7
2 Failures	F2P4	F2P5	F2P6	F2P7

Table 2. Customers and Vehicles per Condition

	nCustomers	nVehicles
Scenario 0: F1P4	12	4
Scenario 1: F1P5	20	5
Scenario 2: F1P6	30	6
Scenario 3: F1P7	42	7
Scenario 4: F2P4	24	8
Scenario 5: F2P5	40	10
Scenario 6: F2P6	60	12
Scenario 7: F2P7	84	14

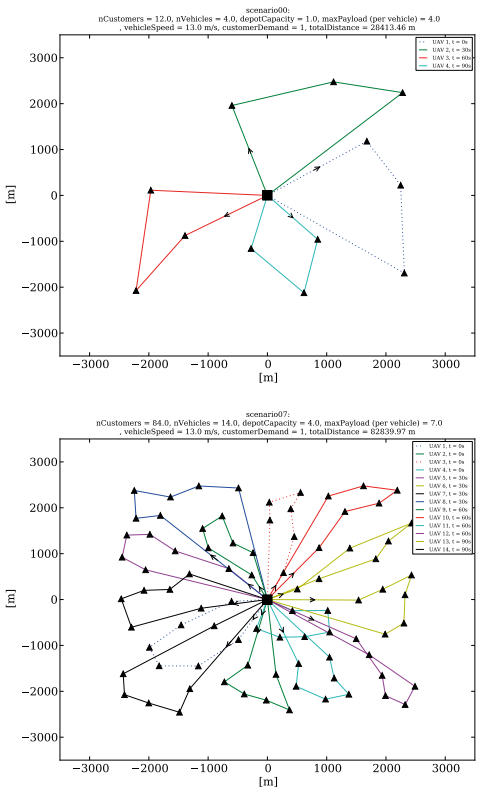

Fig. 3. Two example flight scenarios. The dotted lines indicate the UAVs that failed upon take-off.

Table 3. Control Variables

Variable	Value
Max Flight Time (s)	750
Airspeed (m/s)	13
Service Time (s)	30
Scenario Duration (s)	360
Failure Times (s)	5
Payload Margin (-)	1
Sector Size (m^2)	5000 x 5000
Depot Capacity (-)	30% of nVehicles

All scenarios were created by using an off-line VRP optimization algorithm from Google Optimization Tools, which is a software suite for solving combinatorial optimization problems. Customer locations were generated using a random number generator, where a minimum distance criterion was implemented to prevent location clustering, and the Google toolset was used to calculate the initial optimized plan taking into account the scenario and fixed parameters listed in Table 3.

4.3 Dependent Measures

The emphasis was put on measuring the number of successful missions, the time it took participants to solve the perturbed scenarios and comparing the human solution patterns with the solutions from a static DCVRP algorithm.

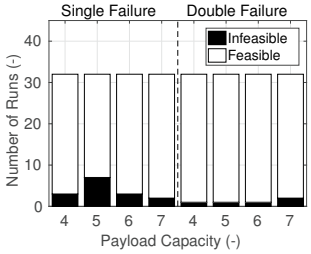

Fig. 4. Total count of unsuccessful missions.

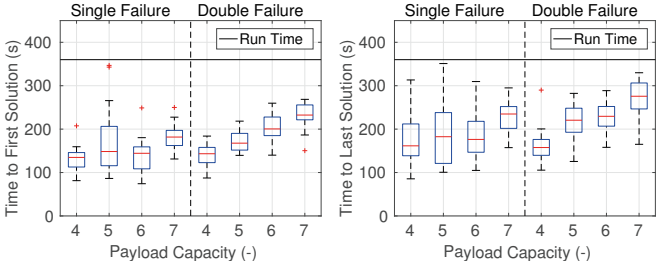

(a) Time to first solution (b) Time to last solution

Fig. 5. Box plot of the time to obtain the first and last solution per condition.

5. RESULTS

5.1 Mission success

Figure 4 shows a bar chart of the number of infeasible and feasible runs per condition. Infeasible runs were counted when participants were unable to meet any of the control goals, i.e., serving all customers, not overrunning the available flight time and not overrunning the depot capacity. The results show that in most cases the participants were able to solve the perturbed scenarios, even the more difficult ones. Surprisingly, the largest number of failed missions occurred in the scenarios featuring a single UAV failure and a relatively low payload capacity, such as condition F1P1 (scenario 0).

5.2 Solution time

Figure 5 shows a box plot of the time to the first and last solution per condition. The solution times are defined as the amount of time between the start of the scenario and the moment the first and last feasible solution was achieved. All participants, having successful missions, could solve the scenarios within the run-time of 360 seconds and on average took about 150 seconds to solve the single failure scenarios and 200 seconds for the double failure scenarios. Observing the trends in the figure, it can be seen that for small problem sizes the time to first solution is relatively constant, whereas it increases for larger problem sizes.

5.3 Solution patterns

Figure 6 depicts the perturbed scenario, the optimized solution (calculated off-line with the Google VRP toolset) and two participant solutions for condition F2P7 (scenario 7). Due to the UAV flight time constraints, the conditions with few customers have a smaller solution space compared to conditions with many customers. Hence, participants

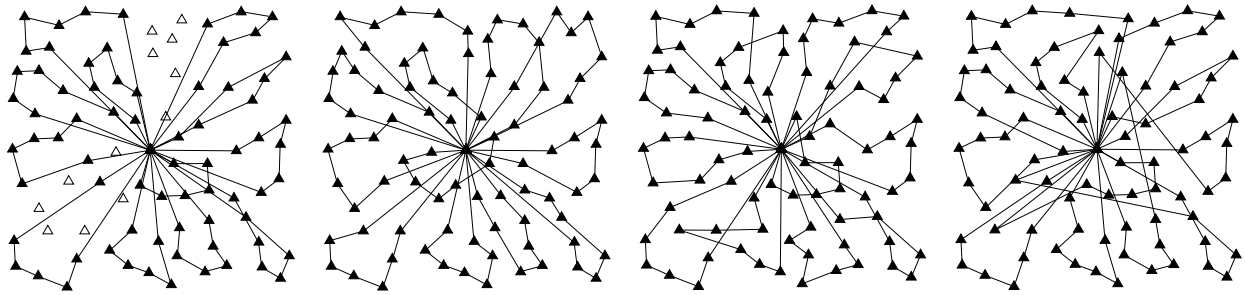

(a) Perturbed scenario (b) Optimized solution (c) Participant 6 solution (d) Participant 1 solution

Fig. 6. Selection of results for condition F2P7.

generally either found one out of a small set of solutions, or were unable to solve the scenario. For this condition, participants with a good strategy were able to come up with solutions that are visually similar to the optimized solution, see Figure 6c, but rarely exactly the same because of the large solution space. Participants with a bad strategy, or participants who focused on satisficing over optimizing generally opted for solutions that visually look more chaotic, as depicted in Figure 6d.

6. CONCLUSION

The goal of this study was to support human control performance in a dynamic multi-UAV DCVRP (featuring vehicle breakdowns) by means of a visual interface. Results show that the developed interface allowed human operators to effectively control perturbed DCVRPs across a range of problem sizes. Interestingly, the participants experienced most difficulty in solving scenarios featuring a single UAV breakdown and less customers. In contrast, computer algorithms generally experience most problems in solving larger problem sizes, as computational time can increase exponentially with the number of customers. For future research, the developed interface could facilitate human-automation cooperation, in which humans can quickly find an initial solution to a perturbed situation, after which an algorithm could further refine and optimize that solution. Further, and a more detailed, data analysis is needed to better compare human solutions with those from a dynamic DCVRP algorithm and explore the limits of human problem-solving activities in more realistic and constrained VRP (e.g., adding time windows to customers and vehicles with different battery capacities).

REFERENCES

Anderson, D., Anderson, E., Lesh, N., Marks, J., Mirtich, B., Ratajczak, D., and Ryall, K. (2000). Human-Guided Simple Search. In *AAAI/IAAI*, 209–216.

Braekers, K., Ramaekers, K., and Van Nieuwenhuyse, I. (2016). The vehicle routing problem: State of the art classification and review. *Computers & Industrial Engineering*, 99, 300–313.

De Wit, J., Tielrooij, M., Borst, C., Van Paassen, M.M., and Mulder, M. (2014). Supporting Runway Planning by Visualizing Capacity Balances of Arriving Aircraft Streams. In *IEEE International Conference on Systems, Man, and Cybernetics*, 3020–3025. San Diego, CA.

Dry, M., Lee, M.D., Vickers, D., and Hughes, P. (2006). Human Performance on Visually Presented Traveling Salesperson Problems with Varying Numbers of Nodes. *The Journal of Problem Solving*, 1(1), 20–32.

Ercan, C. and Gencer, C. (2018). A Decision Support System for Dynamic Heterogeneous Unmanned Aerial System Fleets. *Gazi University Journal of Science*, 3(31), 863–877.

Klomp, R.E., Borst, C., Van Paassen, M.M., and Mulder, M. (2015). Expertise Level, Control Strategies, and Robustness in Future Air Traffic Control Decision Aiding. *IEEE Transactions on Human-Machine Systems*, 46(2), 255–266. doi:10.1109/THMS.2015.2417535.

Laporte, G. (1992). The vehicle routing problem: An overview of exact and approximate algorithms. *European journal of operational research*, 59(3), 345–358.

Laporte, G. (2007). What You Should Know about the Vehicle Routing Problem. *Naval Research Logistics (NRL)*, 54(8), 811–819. doi:10.1002/nav.

Lau, H.C., Sim, M., and Teo, K.M. (2003). Vehicle routing problem with time windows and a limited number of vehicles. *European journal of operational research*, 148(3), 559–569. doi:10.1016/S0377-2217(02)00363-6.

Lin, C., Choy, K.L., Ho, G.T.S., Chung, S.H., and Lam, H.Y. (2014). Survey of Green Vehicle Routing Problem: Past and future trends. *Expert Systms with Appications*, 41(1), 1118–1138. doi:10.1016/j.eswa.2013.07.107.

Macgregor, J.N. and Chu, Y. (2011). Human Performance on the Traveling Salesman and Related Problems. *The Journal of Problem Solving*, 3(2), 1–29.

Pillac, V., Gendreau, M., Guéret, C., and Medaglia, A.L. (2013). A review of dynamic vehicle routing problems. *European Journal of Operational Research*, 225(1), 1–11.

Psaraftis, H.N., Wen, M., and Kontovas, C.A. (2016). Dynamic Vehicle Routing Problems: Three Decades and Counting. *Networks*, 67(1), 3–31. doi:10.1002/net.

Scott, S.D., Lesh, N., and Klau, G.W. (2002). Investigating Human-Computer Optimization. In *Proceedings of the SIGCHI conference on Human factors in computing systems. ACM*, 155–162.

Toth, P. and Daniele, V. (2014). *Vehicle routing: problems, methods, and applications.* Society for Industrial and Applied Mathematics.

Tütüncü, G.Y., Carreto, C.A.C., and Baker, B.M. (2009). A visual interactive approach to classical and mixed vehicle routing problems with backhauls. *Omega*, 37(1), 138–154. doi:10.1016/j.omega.2006.11.001.

Vickers, D., Butavicius, M., Lee, M., and Medvedev, A. (2001). Human performance on visually presented Traveling Salesman problems. *Psychological Research*, 65(1), 34–45.