

Vehicle Routing Problem Approach for Improving Fuel Delivery Scheduling to Austere Test Sites

Sophie Allen, Jack Heller, Gentry Lamb, Kate Leonard, Jared Stewart, and Brian Lemay

Operations Research Program, United States Air Force Academy, USAF Academy, Colorado 80840

Corresponding Author's Email: BLemay@umich.edu

Author Note: The cadet authors are all first-class cadets at the U.S. Air Force Academy, partnering with the 96th Logistics Readiness Squadron. We would like to thank our faculty advisors for their expertise in the completion of this capstone research analysis, as well as the 96th Logistics Readiness Squadron - Eglin Fuels Management Flight for their involvement and diligent contributions to the project. The views expressed herein are those of the authors and do not reflect the position of the United States Air Force Academy, the Department of the Air Force, or the Department of Defense.

Abstract: The 96th Logistics Readiness Squadron at Eglin Air Force Base, Florida conducts an average of 55 million gallons of annual refueling operations across the 464,000-acre Eglin Reservation. Currently, ground refueling operations plans generate suboptimal refueling routes, negatively affecting fuel consumption and operational costs. To improve the efficiency of the refueling operations, we created a data-driven scheduling tool inspired by the Vehicle Routing Problem. The study focuses on 20 key field test sites across the reservation. When applying our scheduling tool to historical refueling operations, we find there are opportunities to significantly reduce mileage on refueling vehicles. The study aims to validate the effectiveness of the proposed tool by comparing suggested routes with current operations and offering actionable recommendations for future fuel distribution logistics. The analysis revealed an average of 6,650 miles saved per year. The scheduling tool's flexibility enables application across diverse installations and operations.

Keywords: Fuel Delivery Optimization, Vehicle Routing Problem (VRP), Logistics

1. Introduction

The 96th Logistics Readiness Squadron (LRS), stationed at Eglin Air Force Base in Florida, manages more than 600 personnel, making it the largest squadron in Air Force Materiel Command (AFMC). They are also home to the Fuels Management Flight, which is responsible for the end-to-end life-cycle logistics over the largest fuel account in AFMC, totaling just over 5 million gallons of useable stock. Currently, they use C-301 refuelers to deliver fuel to austere test sites across the 464,000 acres of the Eglin Reservation. The 96th LRS currently delivers to 70 locations across the reservation, while the Fuels Management Flight services 20 primary locations. Ground refueling operations at Eglin AFB have been critical in successful radar development, missile defense, and munition detonation testing throughout the range. Currently, their fuel distribution scheduling system only considers one order at a time and could therefore be more efficient. For example, the 96th LRS will send one refueling truck to each request that immediately returns to the depot even if it has the fuel capacity to service additional sites. Having to schedule individual trucks and drivers for every request restrains the crew's ability to fulfill other missions. Oftentimes, these austere sites can only be reached by unpaved roads with deep ruts, potholes, and other conditions that affect a truck's ability to travel effectively. Reducing unnecessary mileage is critical to minimizing maintenance concerns. By considering all fuel orders simultaneously when creating a delivery schedule, our team hopes to increase the efficiency of their delivery system. Due to decreasing fuel truck mileage, savings are realized through decreased fuel expenses, maintenance requirements, and work hours required for driving and maintaining the fuel trucks.

1.1 Problem Statement

The 96th LRS delivers more than 55 million gallons of fuel across the Eglin Reservation each year with a fleet of C-301 trucks. Currently, all planning and scheduling is done manually. This manual process is challenging, time-consuming, and expected to result in suboptimal fuel delivery schedules. Thus, our team's objective was to develop a data-driven scheduling tool that reduces the total number of miles driven by the client's fleet of refueling vehicles to enhance the efficiency of diesel fuel delivery to austere test sites. To determine our model's effectiveness, we focus on minimizing the fleet's total distance traveled annually.

1.2 Related Work

The model our team seeks to assist in solving the problem above is similar to a Vehicle Routing Problem (VRP), which has origins more than six decades ago (Dantzig & Ramser, 1959). Dantzig and Ramser dealt with the optimum routing of a fleet of gasoline delivery trucks between a bulk terminal and many service systems. They knew the routes between any two points and the demand for products and used linear programming formulation to obtain a near-optimal solution. Years later, Qian & Eglese (2016) explored methods for looking at routes and schedules for a fleet of delivery vehicles while minimizing fuel emissions. We find this research useful given the environmental constraints for our own problem. They use a column generation-based tabu against real-life data to ensure that it works in a real-life scenario. Beyond only optimizing fuel routes, they center on reducing fuel emissions.

An alternate vehicle routing problem is presented by Doerner, Hartl, and Lucka (2005) through a study of parallelization of the D-Ant algorithm. The authors are primarily concerned with reducing the run time for an Ant Colony Optimization (ACO) algorithm without altering its behavior. They discuss two methods of parallelization, fine-grained and course-grained, and how they are implemented to improve the D-Ant algorithm. They find that by breaking a VRP into sub-problems, they can reduce the time needed to find a solution while still influencing efficiency. Given the scale of our project with the 96th LRS, we find it is not necessary to implement this method of breaking into sub-problems but may be useful for a larger scale project that analyzes all locations the fuels flight serves.

Ichoua, Gendreau, and Potvin (2007) review and classify work in dynamic vehicle routing, spotlighting the idea of uncertainty stemming from new requests. They describe their problem as “a fleet of vehicles in movement to service customer requests that are dynamically revealed over time” (Ichoua et al., 2007, p. 2). They reference many other publications and analyze methods used by previous authors, concentrating on the ability to efficiently integrate new requests into the vehicle’s planned route. This notion of real-time fleet management could greatly benefit our project, for we are working to develop a model that can take parameters that are uncertain as they change daily and produce an optimal schedule and route for the day.

1.3 Organization

In this paper, Section 2 outlines our research approach including the methodology, detailing data collection methods and analytical techniques, thereby establishing the framework for the study. The heart of the paper, Section 3, unveils the core findings and delves into comprehensive data interpretation and pattern recognition. Finally, in Section 4 we encapsulate the study’s significance by summarizing key findings, providing recommendations, and charting potential paths for future research.

2. Data and Methodology

All data gathered for this project was provided directly by the clients at the 96th LRS. The only outside data collection required is generating a distance matrix to quantify the distances between service locations. The data is provided in a spreadsheet and includes a schedule of customer requests for fuel across one year along with a “ground products query,” which includes the dispatched orders and 25 attributes of each order. Attributes that we are interested in are the fuel grade, fuel quantity, delivery location, response time, and fueling time. The 96th LRS also provided our team with a table of the top 20 most visited sites along with their geographic location and coordinates. Although there are 70 locations on the Eglin Reservation, our team focused on the 20 locations we have coordinates for. To maintain operational security standards, all locations have been replaced with generalized names.

We identified general methodologies and developed multiple methods for solving the routing problem. After reviewing current literature in the field of vehicle allocation, we found several methods that align well with our problem statement, such as a Vehicle Routing Problem (VRP), the Orienteering Problem (OP), and Ant Colony Optimization (ACO). Other aspects we consider include a mixed fleet, having two types of fuel service trucks, mixed demand, and non-scheduled delivery requests. While each method has the potential to solve our problem, the VRP closely resembles our specific approach and offers the greatest simplicity of the methods. As such, we employ this method to build a model that accurately embodies the problem we attempt to solve. The main objective of our model is to minimize the total distance traveled by the fleet while meeting demand. While the specific paths between locations will be provided by the client, our goal is to determine the optimal configuration of routes to meet demand. Our final product is a tool that can take daily fuel requirements for each site and build a schedule of truck routes that will satisfy total demand for the day.

3. Mathematical Modeling and Analysis

Using the 20 locations with provided coordinates, our team created a distance and time matrix. The distance matrix represents the miles traveled between each site using the suggested route by Google Maps. Figure 1 below illustrates a condensed form of this distance matrix. The time matrix represents the minutes to travel between each site based on the Google Maps estimation. While we recognize these times can be affected by traffic, we use the average times required to transit from location to location. Future work could include incorporating real-time travel times. With these distance and time matrices, we identify two potential models. The first model, a Daily VRP Solver uses a hybrid genetic search (HGS) algorithm to provide high-quality routes based on the daily information uploaded by the client. The second model, a Pre-Set Route Recommendation, provides an assortment of predetermined routes based on general VRP optimization.

Miles	Depot	Location 1	Location 2	...	Location 17	Location 18	Location 19
Depot	0	25	20.5	...	14.4	18.8	22.2
Location 1	25	0	5.5	...	9.4	3.7	38.1
Location 2	20.5	5.5	0	...	12.3	1.8	41.1
⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮
Location 17	14.4	9.4	12.3	...	0	10.6	31.4
Location 18	18.8	3.7	1.8	...	10.6	0	39.4
Location 19	22.2	38.1	41.1	...	31.4	39.4	0

Figure 1. Distance Matrix

Each model has its benefits and drawbacks. The Daily VRP Solver provides customized, high-quality routes daily, while the Pre-Set Route Recommendation is somewhat less efficient. In that sense, the Daily VRP Solver does a better job limiting the mileage put on the clients’ vehicles and thus is our primary model. The Pre-Set Route Recommendation excels in its simplicity and will be addressed only as an alternative solution. Since it is far less dynamic, it may potentially serve as a “backup” for the 96th LRS. Both models exhibit a simplistic user interface for ease of use. An example interface, shown in Figure 2 below, lets the user select which sites need fuel for the day and the quantity of fuel demanded. Once the aforementioned information has been entered, the model returns the best routes to be used on a given day for refueling operations.

	Required Location	Fuel Amount
Location 1	Yes	120
Location 2	No	
Location 3	Yes	250
Location 4	No	
Location 5	No	
Location 6	No	
Location 7	No	
Location 8	No	
Location 9	Yes	600
Location 10	Yes	400
Location 11	No	
Location 12	No	
Location 13	No	
Location 14	No	
Location 15	No	
Location 16	Yes	75
Location 17	No	
Location 18	No	
Location 19	Yes	400
Main Depot		

Figure 2. User Interface

The VRP can be formulated as an integer programming model, with the objective of minimizing total distance traveled across all routes. To do so, we define the binary variables x_{ijt} to have a value of 1 if the arc from node i to node j is driven by truck t in the solution. We also define the parameter d_{ij} to be the distance from node i to node j . N is the set of all nodes (node 0 is the depot) and T is the set of all trucks. Additionally, D_i is the demand of node i , l_t is the initial load of fuel on truck t , and Q_t is the capacity of truck t . With these definitions, the objective function of minimizing total distance is formulated as in Equation 1.

Below we describe and formulate the constraints our model is subject to. Constraints (2) specifies that each node must be visited exactly once. In our case, a location will never need more fuel than one truck can deliver so we do not need to worry about covering situations where a location needs to be visited more than once in a single day. Constraints (3) guarantees that each vehicle is used at most once. Constraints (4) is a flow balance constraint that stipulates a vehicle arrives and departs from each node it serves. Constraints (5) calculates the total load of fuel delivered by truck t and then Constraints (6) ensures that we do not overload the vehicles by making sure the total demand fulfilled by vehicle t is lower than the vehicle's capacity, Q_t .

$$\text{Min} \sum_{t \in T} \sum_{i \in N} \sum_{j \in N} d_{ij} \cdot x_{ijt} \quad (1)$$

$$\sum_{t \in T} \sum_{i \in N} x_{ijt} = 1 \quad \forall j = 1, \dots, |N| - 1 \quad (2)$$

$$\sum_{j=1}^{|N|-1} x_{0jt} \leq 1 \quad \forall t \in T \quad (3)$$

$$\sum_{i \in N} x_{ijt} - \sum_{i \in N} x_{jit} = 0 \quad \forall j \in N \text{ and } t \in T \quad (4)$$

$$l_t = \sum_{i \in N} \sum_{j=1}^{|N|-1} D_j \cdot x_{ijt} \quad \forall t \in T \quad (5)$$

$$l_t \leq Q_t \quad \forall t \in T \quad (6)$$

3.2 Analysis

To evaluate the effectiveness of our model, we compare how the total distance traveled with the recommended routes differs from that of current operations. We simulate one day of operations by using historical daily schedules and demands to test our model, and calculating the total mileage required to fulfill all daily fuel requests.

First, we analyze the current operations. If we simulate using demand at Locations 1 through 7, the vehicle fleet would travel 212 miles and require 6 trucks to make those deliveries. When we plug those same locations and demand into our model, we discover that the total fleet distance traveled would be 118 miles using only 2 refueling trucks. While there is still room for improvement, our model performs better than current operations, saving the 96th LRS nearly 100 vehicle miles while dispatching fewer trucks. Figure 3 below illustrates the model output and how the routes are recommended. In the table, each row represents a different vehicle while the columns represent the stops that vehicle makes along its route. Also note that all vehicles must return to the depot to complete their route in full. Currently, the Fuels Service Center rarely has more than 3-4 of these locations request fuel in a day. However, with the discoveries we have made, we hope to inform decisions of the leadership to allow for these requests to “build up” to maximize the efficiency of the fleet. We discuss this more thoroughly in the recommendations section below.

	Stop 1	Stop 2	Stop 3	Stop 4	Stop 5
Truck 1	Main Depot	Location 3	Location 1	Location 2	Main Depot
Truck 2	Main Depot	Location 6	Location 5	Location 4	Main Depot

Figure 3. Route Recommendation Model Output

In addition, our team performed rudimentary simulation exercises to compare longer-term effects of the model. Since the raw data that was provided contained errors that made it impossible to calculate exact demand distributions for each location, we simulated the demand based on the data known to be correct. Our simulation ran over a one-year period. Each day, anywhere between one and five locations were randomly selected and assigned a random demand for the day. We would then calculate the distance traveled under current operations, out and back to every location, as well as the distance traveled with our model. The model outperformed the current operations on nearly every occasion, with a handful of discrepancies that occurred when high demands were observed at all the randomized locations for the day. Over the course of a year, we discovered that our model could, on average, improve total fleet mileage by a respectable 14-20%. While this simulation does not perfectly represent the fueling operations of 96th LRS, it informs us of the potential savings that can be achieved with our model. Although our team was unable to retrieve accurate data on operating costs, maintenance costs, and acquisitions costs to calculate total costs, estimates indicate that we could save on the order of thousands of dollars a year. Future work could dive deeper into calculating cost savings and enhancing simulation methods.

4. Conclusions, Recommendations, and Future Research

The current fuel delivery system for the 96th LRS operates below optimal efficiency, resulting in excess mileage on vehicles, which incurs additional operating and maintenance costs. Employing a Vehicle Routing Problem model, with an objective of minimizing total fleet mileage, we identify daily optimal routes, leading to significant savings in distance and costs for the 96th LRS. Further savings in terms of personnel hours are significant, as our model takes less than a minute to run, start-to-finish. Based on historical data and foundational simulation, we anticipate a baseline improvement of 16%, effectively reducing annual mileage by more than six thousand miles compared to the conventional out-and-back approach. While we lack the necessary data to calculate savings in other areas, such as operational, maintenance, and acquisition expenses, our approach significantly improves operational efficiency and cost-effectiveness for the 96th LRS.

4.1 Recommendations

Moving forward, we propose several recommendations to maximize the efficiency of refueling operations. Foremost, we recommend an adjustment in current ground refueling scheduling. Specifically, we endorse a new process that sets specific, limited dates for ground deliveries. This may resemble traditional trash disposal operations, where trucks are scheduled to come on specific days of the week, so ground deliveries would only be made once or twice a week. By limiting the days on which ground deliveries are conducted, the Fuels Management Flight can better prepare for meeting customer demand. If they know ahead of time which days will require ground deliveries, they will have more flexibility in scheduling drivers to meet other mission requirements such as aircraft refueling. This does come with a tradeoff in the sense that the locations needing the fuel might not be able to get the fuel they need in time, depending on how often they require refills. Additional analysis of historical data could prove the feasibility of this recommendation but requires better data collection processes.

Additionally, our tool becomes more efficient as more customers are added to the problem. For example, when there are only three customers to deliver to, mileage savings are minimal in the 5 to 10 mile range. When eight or more customers are added, mileage savings increase drastically into the 20 to 100 mile range. If the 96th LRS can adjust their current scheduling processes to allow for planned delivery days, they will be able to optimize fleet operations by lowering costs and saving thousands of miles per year on vehicle wear. Finally, we recommend enhancing data collection by standardizing the input process for all locations to support data analysis efforts. This entails implementing a system that ensures uniformity in location entries, mitigating discrepancies that arise from multiple entries from varying users who capture location names and other attributes differently. Something as simple as selecting from a drop-down list of predetermined values would alleviate this issue.

4.2 Future Research

Our scheduling tool holds significant promise for further development and refinement. In the context of the 96th Logistics Readiness Squadron at Eglin Air Force Base, potential avenues for future research include fine-tuning the algorithm to account for dynamic factors such as traffic patterns and weather conditions and integrating real-time data for more accurate route planning. Additionally, exploring the tool's scalability to other Air Force bases, deployed locations, or during emergency response situations presents an opportunity for extending its utility and impact across a broader operational context. Further research efforts may focus on adapting the tool to accommodate varying logistical infrastructures and operational requirements unique to different bases.

5. References

- Dantzig, G. B., & Ramser, J. H. (1959). The truck dispatching problem. *Management science*, 6(1), 80-91.
- Doerner, K. F., Hartl, R. F., & Lucka, M. (2005). A parallel version of the d-ant algorithm for the vehicle routing problem. *Parallel Numerics*, 5, 109-118.
- Ichoua, S., Gendreau, M., & Potvin, J. Y. (2007). Planned route optimization for real-time vehicle routing. *Dynamic Fleet Management: Concepts, Systems, Algorithms & Case Studies*, 1-18.
- Qian, J., & Eglese, R. (2016). Fuel emissions optimization in vehicle routing problems with time-varying speeds. *European Journal of Operational Research*, 248(3), 840-848.