A Metaheuristic Approach to Optimizing a Truck and Drone Delivery System

By: Rick (Yue) Kuang
Advisor: Dr. Mohammad Moshref-Javadi

Topic Areas: Last Mile, Urban Logistics, Artificial Intelligence

Summary: This research contributes to furthering the understanding of the application of autonomous flying drones for improving parcel delivery performance within the constraint of the current state of the technology. An optimization algorithm is created to obtain the best route of truck and drones for delivery of packages. The analysis completed and lessons learned can be broadly applied to a wide range of truck-and-drone routing problems. The software developed from this research has additional value as blueprint for a future product that can provide tactical scheduling and decision optimization functionality to drone embedded systems in the field.

Before coming to MIT, Rick worked as an internal strategy consultant for AT&T as well as held positions in military, consulting, and technology companies. He holds B.S. and M.S. degrees in electrical engineering, a B.A. in economics, and a MBA. He currently holds the rank of Major in the USAFR.

KEY INSIGHTS

1. **Aerial drones can improve last-mile delivery performance over truck alone.**
2. **A genetic algorithm can be used to find and optimize last-mile routing solutions.**
3. **The efficient launch and retrieval of drones may determine success of drone deployment in dense urban centers.**

Introduction

The last-mile delivery problem from a DC to a group of customers can be accomplished in three ways with respect to the usage of drones:

1. The deliveries can continue in the traditional method with just trucks running from the DC to each end customer without the assistance of drones.
2. The deliveries can be accomplished using only drones ferrying packages directly from DC to each customer.
3. The deliveries can be accomplished with trucks that take packages from DC and then deliver them to customers by a combination of truck deliveries and/or deploying drones from the truck.

Method (2) is difficult for generalized deployments due to the number of deliveries that will require going back-and-forth from full length of the distance of DC to each customer, as well as the inability for drones to carry large packages. It can be a supplementary solution for a niche percentage of deliverables consisting of high value parcels that are light and going to places inefficient to route the truck. Method (3) is a widely accepted method. For method (3), important considerations have to be made on where and when to deploy the drone, as well as how many drones should be available for each truck route. If these considerations are not made well then the solution of (3) may be suboptimal to (1); squandering the drone use-case and underselling the potentials of the technology.

This research explores the optimization of a specific Mixed Integer Linear Programing (MILP) mathematical formulation for a method (3) truck-and-drone paired delivery system. In this system, a truck leaves the DC carrying multiple aerial drones. At any given delivery point, the truck can choose to launch one or multiple drones it is carrying to make single delivery trips to selected customers and thereby saving the truck from making that trip. The drones can then return to the truck as it moves along its more optimized route and the trip ends with the truck returning to the DC with the all drones it left with. When optimized for total costs of the delivery given real world cost and performance assumptions, the investment in drones can be shown to have significant savings over traditional usage of trucks alone in some conditions. The brute force optimization for a problem as complex as the one formulated in this MILP is highly cumbersome and it takes long computation times to obtain an optimal solution with more than 10 customers. In a real world scenario, where a truck may deliver to as many as 100 customers on its route, it is difficult to assert that a method that works well with just 10 customers can sustain that advantage without additional analysis. In order to justify the employment of this synchronized truck-and-drone delivery in practice, it is important to reduce the optimization time for real world scale problems to minutes from possibly days or weeks.

There are generally three paths to deal with the real-world optimization complexity of problems like this. One path is focused on algorithms that reduce the complexity of a realistically-scaled problem to an extent that its exponential scaling can be tolerated. As the network becomes much larger, a different set of paths are necessary to achieve efficiency. This set of paths are termed “heuristics” where inexact solutions are yielded but the algorithm complexity is reduced from exponential to polynomial. Under the general heuristics umbrella, one path is to focus on short-cut methods specific to routing; another path focuses on the transformation of the problem space to fit a discipline-agnostic “metaheuristic” method. Metaheuristic methods typically have a record of success across many disciplines.

This research uses metaheuristics to optimize a very specific type of the truck-and-drone delivery problem for large scale networks.

Approach

The classic TSP, its derivative VRP, and the innovative truck-and-drone problems are universally accepted as “NP hard” as of the writing of this thesis. The effort required to solve these problems scales exponentially rather than as a polynomial with respect to the number of nodes in the problem space. An algorithm towards an exact solution to this type of problem can be computationally inefficient for a large network. Every unique sequence of order of delivery is a unique solution. When the problem is explored with drones, there is a large range of combinations of which customers are served by drones and which customers are served by trucks for each sequence of delivery order. The problem this research explores introduces additional complexities, including how many drones a truck should carry, where to launch a drone, how many drones to launch, and where to pick up used drones.

The complexity of the problem and the real-world applications of its solution made it a good fit for heuristics problem solving. In assessing strategic impacts of brand new techniques, the exact optimal solution should not be much more informative than one a few percentage points away. The magnitude of complexities of the problem makes it more ideal to leverage higher-order metaheuristic artificial intelligence techniques rather than simple rule-based heuristics. This research uses metaheuristics in an algorithm to produce a “good enough” solution to the truck-and-drone problem. The algorithm orchestrates pseudo-random generations of many feasible solutions that are selected into a solution set. The quality of the solution set becomes better and better over time until continuous improvement becomes so difficult that there is confidence that a good solution has been reached.

Algorithm

The base framework of the algorithm is the genetic algorithm (GA). The principles of GAs adhere closely to their biological roots. Parents pass genetic information to their children in form of information captured in chromosomes. At each generation, the parental traits that make offspring most suited to survive are pass on. Those offspring stuck with less desirable traits fail to reproduce. Generation after generation, the genetic material gets stronger without the benefit of specific guidance.

At the beginning of each generation, the solution generation process builds up a large population of solutions using either two strong parental solutions from the previous generation, mutating one parental solution from the previous generation, or using pseudo-random methods with no parental information. The initial generation of solutions use only pseudo-random methods. As generations progress, more solutions in the base solution set are dependent on parental solution guidance. At the end of each generation, the algorithm updates machine learning paths data and strengthen the pseudo-random path selection that is core to the solution generation process and then prunes the solution set to retain the most fit parent solutions for the next generation. The process continues to the next generation until the exit conditions are met. The exit conditions are that a minimum amount of execution time has been exerted in the effort and also that a set amount of generations or processing time has passed without a meaningful improvement to the solution quality.

All algorithms were developed using the Spyder 3.3.3 Scientific Python Development Environment on Python 3.7. and implemented on a PC desktop with Intel Core i7-7700 3.60 GHz processor and 16 GB of RAM running on Windows 10 operating system. Solutions were captured in excel and explored with desktop version of Tableau 2018.3.

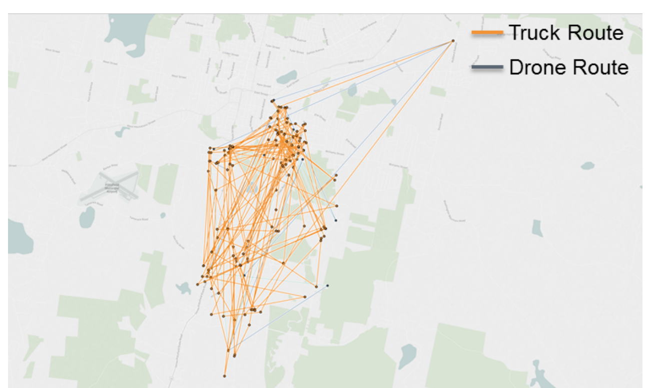

Figure 1: A poor random solution before optimization

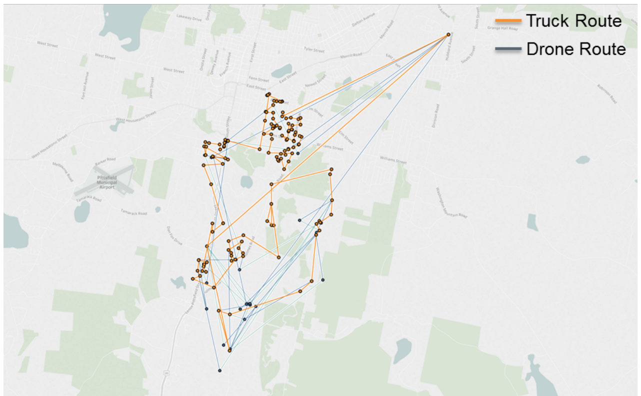

Figure 2: An optimized solution returned by algorithm

Results

This thesis took a metaheuristic approach that played the game of generating pseudorandom solutions round after round that fit within the feasibility constraints of the real world application of the problem. At the end of each round, the algorithm learns which decisions tend to be associated with high quality solutions and bias the next generation of solutions towards generating better solutions. This is not unlike what powers the concept of survival of the fittest in genetic engineering.

A truck capable of efficiently deploying and retrieving autonomous drones to make delivery will generally be able to see a savings in the form of reduced time cost to the delivery truck dampened by some costs to acquire and operate the drones. While having more drones available, having drones that can travel to targets with Euclidean distance, having faster drones, and having drones with more battery range can all have positive effects on the cost savings of the system, the effects are not equal.

The primary improvement opportunity for the studied delivery system is to have drones that are capable of traveling in a straight-line Euclidean distance towards their target rather than following the road systems. In rural areas, this is more likely. If delivery is made in a very dense metropolitan area, it is unlikely to be able to circumnavigate buildings in a direct line. However, drone delivery technology should be able to take advantage of more freedom of movement from ground traffic and set road infrastructure to more efficiently get to its target than trucks. Investment should be made for autonomous drones to handle building obstruction and in-air congestion so as to able to take advantage of the Euclidean distance savings. For rural maps tested, the baseline optimal cost estimate of a single-truck solution is improved by an average of 7% to 9% when 2 drones are carried and these drones travel by actual road distances between customer points. The improvement is at an average of 9% to 12% when the drones are capable of flying directly and the distance becomes Euclidian. This is roughly what can be achieved by adding 2 extra drones without incurring the added asset costs and planning complexity.

The next largest area of opportunity is increasing the average speed of the drone. Improving drone technology to specialize towards high speed delivery of very light packages tends to greatly improve the cost savings of using drones. This savings comes from being able to reliably send drones off to make delivery and come back quickly enough to avoid incurring wait penalty from the truck at pickup points. There is an interesting tradeoff that takes place. Sending fast drones to faraway places saves the truck the cost of making that trip, but that drone will be unavailable for use until much later when the truck makes its way to a pickup point. However, if drones are sent to make short delivery trips, the savings attained from those drone trips may not be worth the two minutes of overhead cost the truck needs to launch and then retrieve the drone.

For map instances that were tested in the experiments, adding drones to a truck has diminishing returns after the second drone. It would appear that a truck cannot do much better with 4 drones in its carrying capacity than it can with 2 or 3. Going with higher than 4 drones may be not worth the cost of acquiring additional drones. For the map instances that were tested, a 30-minute range appears to be sufficient to gain most of the savings opportunities. Adding more battery capacity and drone range does not predictably increase cost efficiency for the system under study.

Finally, while both rural maps of low density of demand and urban maps of high density of demand are studied, the conclusions are mostly attained from test results with rural instances. For the parameters that were used, urban instance maps where the distances between any two given customers are only around 3 kms indicates little opportunity for drone use. This is because the assumptions are that the overhead for deploying and picking up a drone is 2 minutes of waiting time incurred by the truck. A truck could have traveled 1.3 kms at 40 km/h. Much of the efficiency advantage the drone can bring over such short distances of travel between customer nodes would be negated by the truck overhead penalty. If drone launch and retrieval technology were to be improved to a point where it no longer imposed overhead cost to the truck, then drone technology should be able to bring dramatic savings opportunities to the dense urban customers as well.

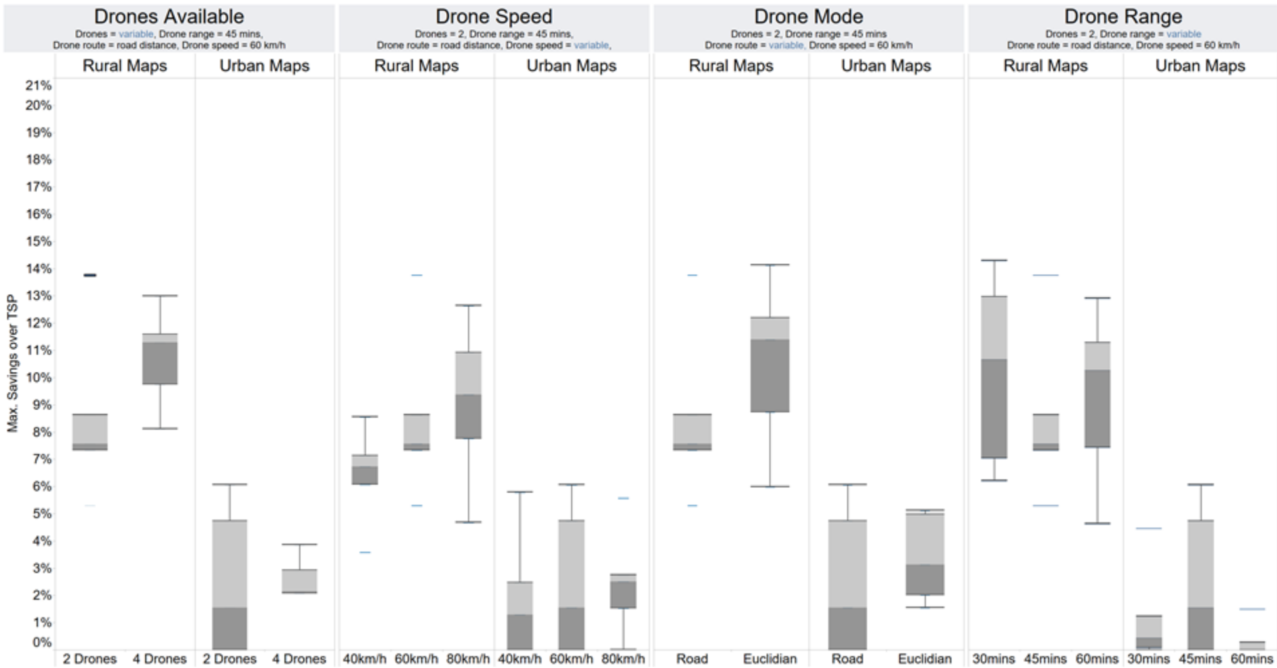

Figure 3: A comparison of sensitivity tests

Conclusions

There are many challenges and complications associated with finding the optimal deployment of a multimodal truck and drone delivery system for last mile delivery. While computing power limits how large a problem can be solved exactly, numerous methods can be used to quickly and efficiently arrive at a good solution that is for all practical purposes almost as good as the exact optimal solution.
After many added features and algorithm tuning, this generalized method was able to produce high quality solutions within a tolerable margin. This algorithm was able to produce these quality solutions with great consistency and within a short period of computation time across a number of maps of real demand data from MIT MLL.

Future work expanding on the approach of this research should concentrate on either applying the algorithm created in this thesis to more realistic and robust problem constructions, or finding new approaches to optimize the problem as constructed. No matter which avenue of improvement future researchers pursue, they can use the algorithm framework developed in this thesis to quickly test new variables and parameters. The algorithm itself is also very adaptable, able to incorporate other methodologies into the overall Genetic Algorithm structure. This research made it a little easier for drone technologists and last mile strategists to assess their risks and opportunities and make the right decisions on where and how they can improve or deploy truck-and-drone systems to enhance the lives of their customers.