ElasticRoute

Hybrid Heuristics for High Speed Route Optimization

Table of Contents

Preface

Last-mile delivery has always been expensive. It accounts for close to 30% of the total cost of the supply chain required to move goods from the manufacturer to the final recipient. High manpower costs in the planning and facilitation of these deliveries further push up the cost of last-mile deliveries. According to a survey report by Retail TouchPoint in 2018[1], the biggest challenge in the last mile for retailers is customers' expectations for fast and cheap deliveries.

The challenges faced by the last-mile logistics industry is exacerbated by the e-commerce boom that is creating massive volume for last-mile delivery. While there exist tools and systems designed to alleviate these problems, there remains a pertinent challenge that is largely still unresolved. That is, being able to efficiently generate optimized delivery routes for all the drivers in time for dispatching. Speed is critical in today's last-mile deliveries. The ability to plan the routes for hundreds and thousands, sometimes even in the tens of thousands of deliveries is extremely challenging. Besides this bottleneck faced by logistics providers, the need to be able to adjust the planned routes presents an even greater challenge. Last-minute changes of delivery addresses, dates, time windows, and last-minute adding of stops are all part and parcel of last-mile delivery operations.

In this paper, we will discuss the challenges faced in route planning and how existing algorithms are largely insufficient. Finally, we will present how the use of a hybrid heuristic method derived from the combination of straight line heuristics and real road distance metrics presents the best solution to solving the challenge of planning efficient routes at high speed for last-mile deliveries today.

[1] Retail TouchPoint (2018), Survey Report – Are You Winning The Last Mile Of Retail?

Introduction to the Problem

Route Planning in last-mile delivery is a difficult problem to solve. Increasingly, many suppliers and distributors no longer maintain in-house fleets, and instead rely on third-party logistics (3PL's) for last-mile deliveries from the supplier's depot or warehouse.

3PL couriers have it especially hard because of the lack of time between when they receive their delivery manifests from suppliers and retailers for the following day and when they need to start dispatching their drivers for deliveries. These couriers often receive their manifests less than 24 hours before their dispatch, or in extreme cases, less than 12 hours. 3PL's then have to come up with an optimal route plan within this short period of time.

This is further exacerbated by the emerging trend of next-day or even same-day delivery with retailers trying to push the limits as to when they can submit their manifests to the 3PL's. Retailers often offer "same-day delivery before 12pm" and "next-day delivery before 5pm" options. It is to the retailer's advantage if they could allow the customer to confirm their order later and still opt for same-day or next-day delivery. However, this poses a great burden on the 3PL's as they will then receive their manifests at a later time. Accounts of planners working overtime through the night to plan the routes for the next day are not unheard of. Rushing through the planning process often produces suboptimal routes, which leads to other problems such as missed time windows on the customers' end.

This is why speed is key in last-mile delivery. Couriers need to be able to quickly plan optimized routes for their drivers regardless of how late they receive their manifests. Route planning softwares that can quickly compute optimal routes will thus prove to be a valuable asset to the company by saving a significant number of man-hours from manual route planning. Fast route planning softwares will also give these courier companies a competitive edge by offering, to their retailer and supplier clients, a faster turnover rate between receiving the manifest and dispatching the drivers, who in turn can offer their end consumers a later cut-off time for same-day or next-day delivery.

In this white paper, we analyze the mathematical and computational problems used to model the process of last-mile delivery and dig into the reasons why they are so difficult and expensive to solve by traditional Route Planning and Route Optimization softwares. We explore the possibilities of using Hybrid Heuristics to speed up the route planning process and investigate its performance. We also introduce our new routing service, ElasticRoute, and explain how it could empower courier companies to be agile with their route planning and vehicle dispatching.

Routing Problems in Last-Mile Delivery

Route Planning is the process of assigning your fleet of vehicles to serve a number of stops in an order that minimizes fuel costs and travel time. Many mathematical problems have been formulated in an attempt to model and solve the route planning problems in Last-Mile Delivery. The Traveling Salesman Problem (TSP) and the Vehicle Routing Problem (VRP) are the most common problems used to model Last-Mile Delivery. As NP-hard problems, it becomes excruciatingly difficult to derive a solution in a reasonable amount of time as the time complexity grows exponentially with an increase of number of stops and vehicles.

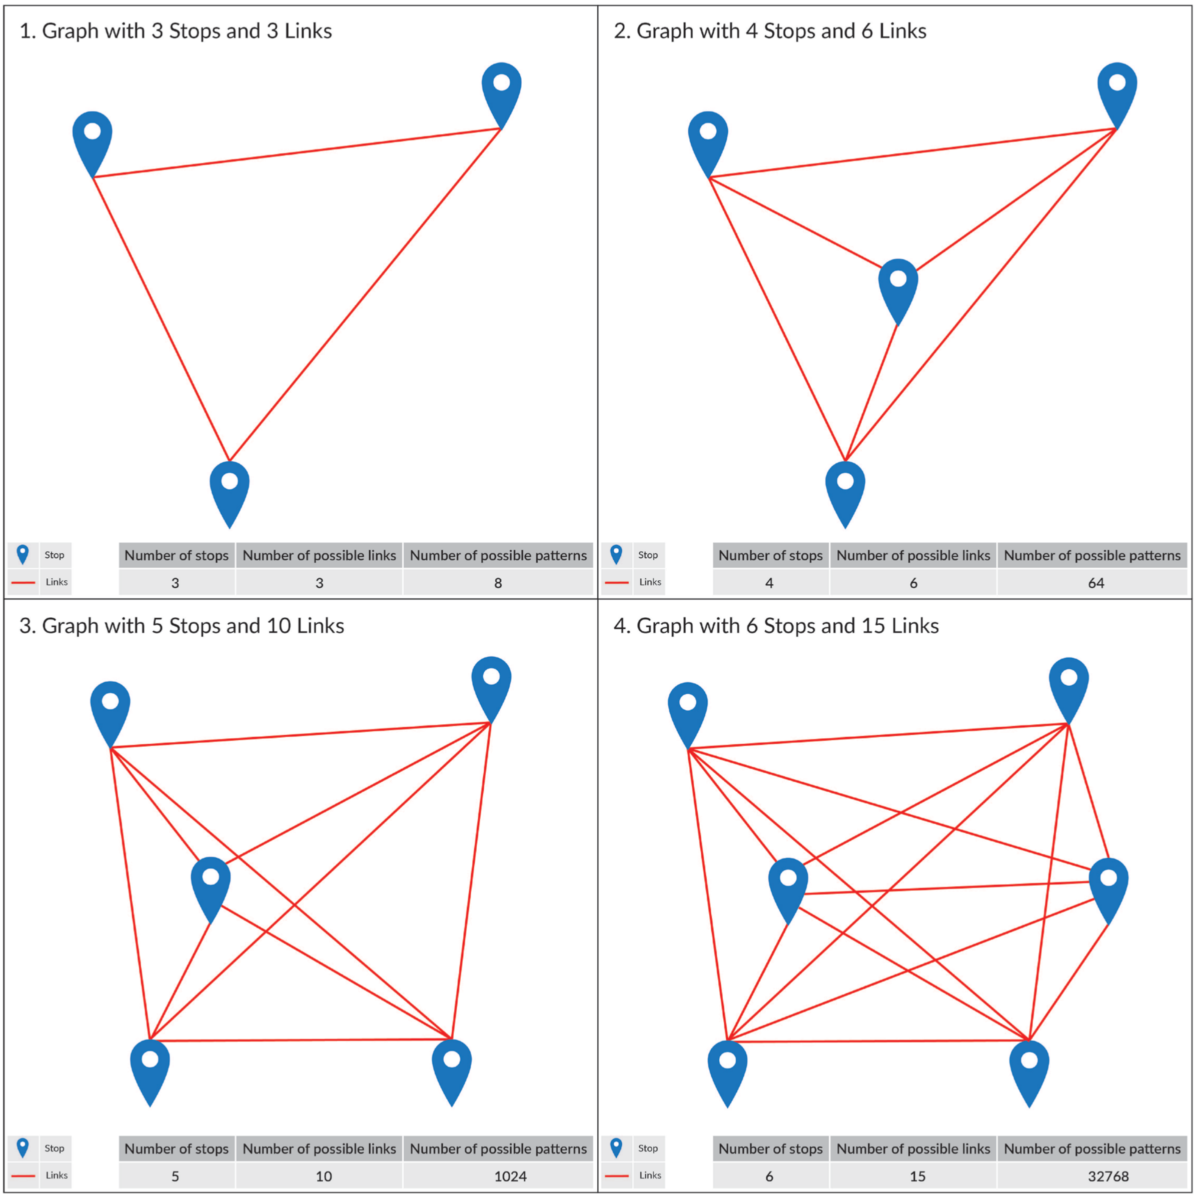

Figure 1: Illustration of the NP-hard problem of exponential increase in complexity with increase in data size

Figure 1 shows how complex the number of paths grows with just a few stops. However, this is far from the scale of what couriers in Last-Mile Delivery face. The table below shows the total number of combinations (“patterns”) each driver can take for a different number of stops, and the number of such patterns that are valid solutions to the corresponding Travelling Salesman Problem (TSP).

Number of Stops	Number of Possible Links	Number of Possible Patterns	Number of TSP Solutions
N	N(N-1)/2	2^L	(N - 1)!/2
3	3	8	1
4	6	64	3
5	10	1024	12
7	21	2097152	360
8	28	2.68E+08	2520
9	36	6.87E+10	20160
10	45	3.52E+13	181440
100	4950	1.24E+1490	4.67E+155

Table 1: Complexities of the TSP

As complicated as it gets, solving the Travelling Salesman Problem and the Vehicle Routing Problem is key for route planners to optimize their drivers’ routes. We look at how route planners use the TSP and the VRP to model Last-Mile Delivery and their proposed solutions.

The Traveling Salesman Problem

The well-known problem in logistics and operations research: given a list of cities and the distances between each pair of cities, what is the shortest possible route that visits each city and returns to the original city?

What is it used for?

In the context of last-mile delivery, the TSP is often used as a simplified model of the VRP. Route planners pre-assign stops to vehicles before coming up with a route for each vehicle. For small TSP's, say 20 stops per vehicle, using human intuition in solving the TSP by hand may produce acceptable results, but TSP's larger than that would often result in inefficient routes. Route planners then defer the work to route optimization softwares that use several techniques and heuristics to try and produce an optimal route faster than human planners.

Why is it difficult to solve?

The TSP gained its infamous reputation in Operations Research because of the sheer number of available solutions to explore and evaluate. The number of permutations grow exponentially – a TSP with 7 stops has 360 solutions, while for a TSP with 10 stops the number of solutions explodes to 181,440[2]. With larger scale TSP's on the magnitude of thousands or tens of thousands of stops, the number of solutions is so large, the number would have so many digits it wouldn't fit into this page (or possibly longer than this entire paper).

There are significant cost factors in finding solutions for the TSP as well. Many Operations Research Tools that are able to solve the TSP (and in extension, the VRP) require a fully populated distance matrix in order to derive accurate solutions. The distance matrix is a table of distances for every origin-destination pair in the problem. However, using this distance matrix can get expensive really quickly. Third party map services such as the Google Maps platform charge a substantial amount for calculating road distances between two points. For a theoretical problem of 1000 stops, the distance matrix would contain one million elements. The price for evaluating this huge distance matrix would then cost 4100[3] USD. Daily operational costs of this scale are a luxury that only the biggest of conglomerates could afford. These costs presents a high barrier to entry to smaller enterprises looking to enter the industry.

[2] The number of possible solutions to a TSP with n stops is (n-1)!/2.
[3] As of September 2019, Google Maps' Basic Distance Matrix API lists its pricing at 5 USD per 1000 distance matrix elements, before applying volume discounts. More information can be found at the Google Developers' page: https://developers.google.com/maps/documentation/distance-matrix/usage-and-billing

Limitations of the Traveling Salesman Problem

As the elephant in the room of Operations Research, decades have been poured into finding solutions to the problem. While the TSP is an interesting problem to study, it has several limitations that prevent it from modelling route planning in Last-Mile Delivery effectively. We investigate the real-world implications of these limitations based on accounts from logistics companies who have shared with us their workflows of route-planning using TSP-like methods and strategies.

What vehicle serves this stop?

This would be the question that most route planners ask when trying to use the TSP to model the route planning problem. In order to set up the TSP, route planners have to first assign a list of stops to each vehicle, then solve the TSP for each vehicle where each individual vehicle is the “traveling salesman”. However, the task of assigning vehicles to stops is a problem on its own that needs to be optimized, as poor assignment can lead to poorly formulated TSP solutions that in turn leads to inefficient routes for drivers.

Problems with Manual Geographical Zoning

A common answer to the “what vehicle serves this stop” question is to make use of geographical zones. Manual route planning workflows and route planning softwares often require route planners to first divide the city (or country) into several geographical zones, and then assign one or more drivers to each geographical zone. However, this method is not scalable and is prone to errors. If not drawn properly, human-demarcated zones can leave certain areas outside the coverage of any zone, causing the routing algorithm to mark the stop as unserved. (Figure 2) It would require significant amount of manual labor to demarcate these zones without error.

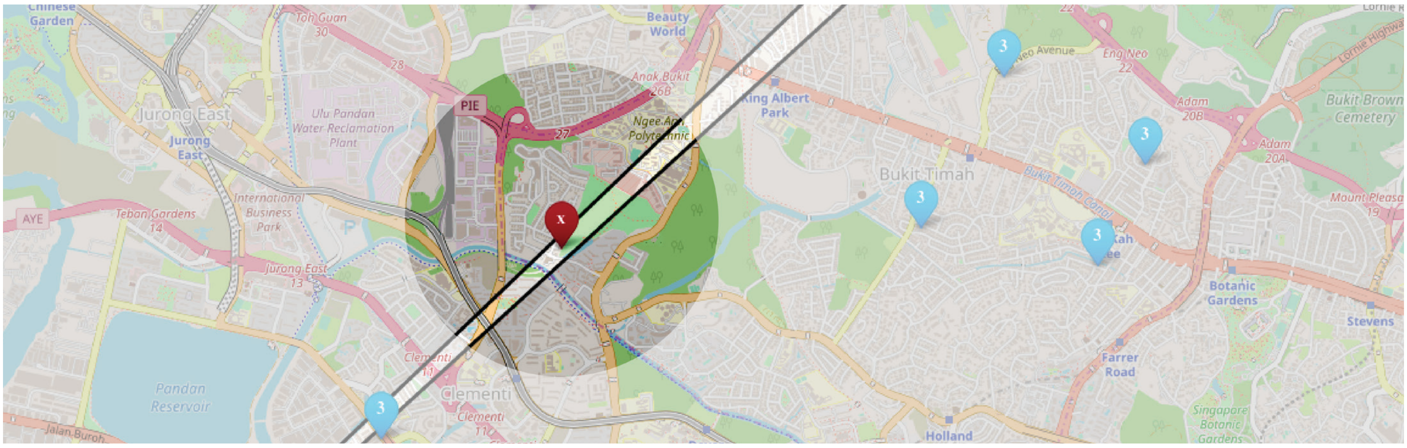

Figure 2: An unserved stop due to hand drawn zones not exhaustively covering all land areas

The strategy of assigning vehicles to stops via geographical zones only works well if there is an equal distribution of stops across the different zones. In addition, fixed geographical zones cannot react well to spontaneous surges and variations in demand across these zones. The distribution of stops across zones may vary on a day to day basis due to many unpredictable factors such as consumer demand. A collection of geographical zones drawn today might not respond well to demand the next day, causing some geographical zones to contain much more stops than the others. In this case, TSP algorithms will mark these excess stops as unserved (Figure 3). Manual route planners may react to this by adding more vehicles to the fleet to serve that region, but this unncessarily increases operation costs. These surges in demand can be accommodated by strategically redrawing the zones to cope with the demand, but route planners cannot afford to redraw these geographical zones every day.

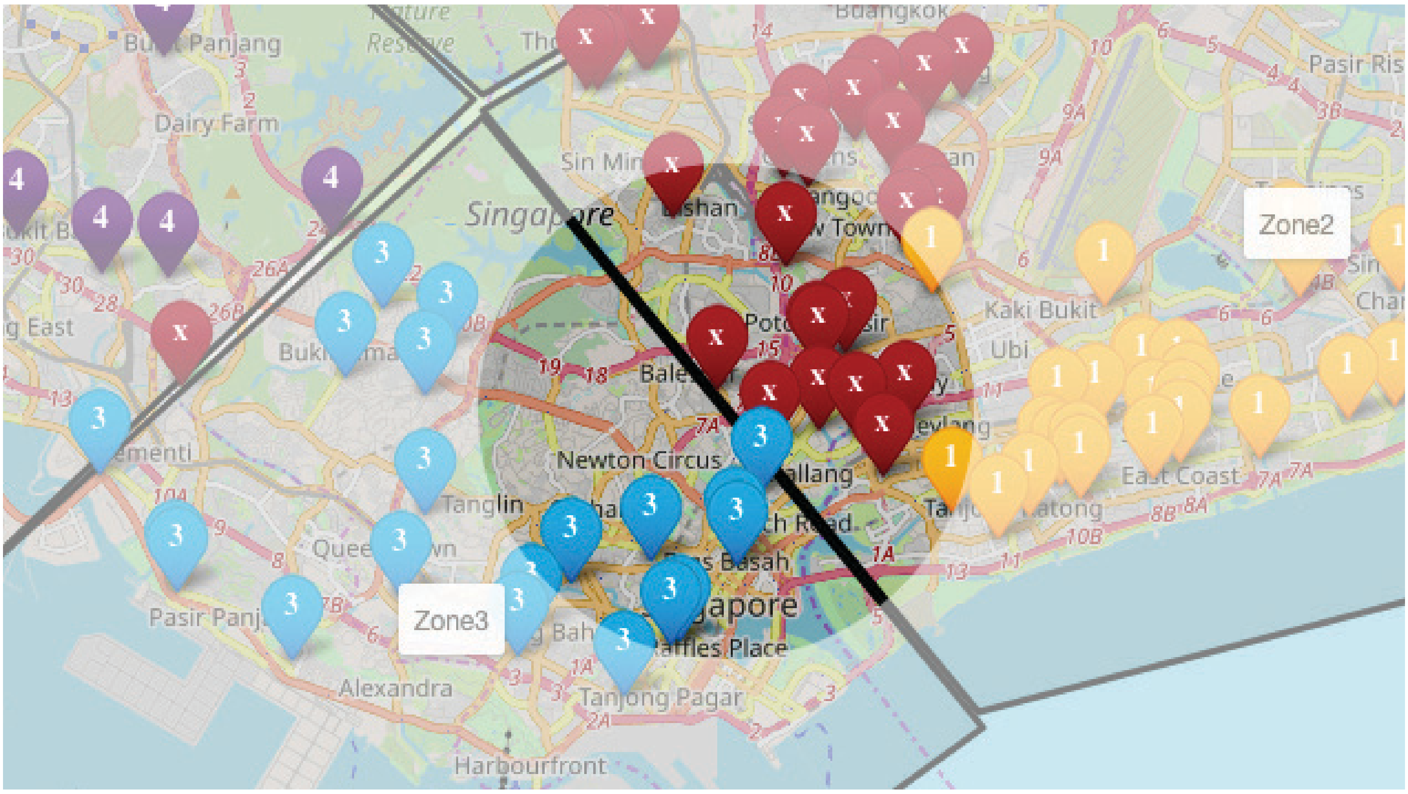

Figure 3: Screenshot showing the large number of stops marked unserved (in red) due to an overallocation of stops in Zone 2.

It can be clearly seen that just using the TSP to model route planning in Last-Mile Delivery is inadequate, and couriers should avoid using solutions that can only solve TSP's if they are looking to optimize their Last-Mile Deliveries. To accurately represent the problems of route planning in Last-Mile delivery, we need to refer to the Vehicle Routing Problem, which is a generalization of the Traveling Salesman Problem that looks to provide answers to questions like the allocation of vehicles to stops.

The Vehicle Routing Problem

The bigger elephant in the room of Operations Research, the Vehicle Routing Problem (VRP) now asks the question: how should vehicles be allocated to stops in a way that optimizes traveling time? As a generalization of the TSP, the VRP is much more difficult to solve in large numbers due to its inherently larger solution space. We look at how being able to quickly solve the VRP on demand generates value for couriers in the field of Last-Mile Delivery.

Value of the Vehicle Routing Problem

Because the VRP helps Route Planners to assign vehicles to stops in the most optimal fashion, the VRP is among the most reliable models to aid route planners in routing their fleets. In addition, subproblems of the VRP models situations that are often faced by couriers in the context of Last-Mile delivery, such as having to consider load and volume capacities of vehicles when planning routes and accounting for the customer’s chosen time windows. Being able to solve the Vehicle Routing Problem with these additional constraints will save on re-plans and aborted runs due to violations of load constraints and improve customer satisfaction by fulfilling the time windows. The fulfilling of time windows is especially important for retailers – from our existing Detrack Proof of Delivery system, we see that a missed time window will almost definitely (96%) end in a negative review from the customer. Solving the VRP is hence very valuable to courier companies in the Last-Mile Delivery field.

Challenges in solving the Vehicle Routing Problem

We had explored why the Traveling Salesman Problem is already a difficult problem on its own. The guesswork of allocating the right vehicles to the right stops, combined with the additional constraints such as loads and time windows, makes the VRP a much more challenging problem to solve than the TSP. Existing VRP solutions also have caveats that prevent them from being deployed in the everyday workflow of couriers in Last-Mile Delivery. In particular, VRP solutions often run into the “choose two out of three” problem: high speed, natural routes, and low computation costs. Existing VRP solutions and heuristics tend to only offer at most two out of the three above critical points to couriers.

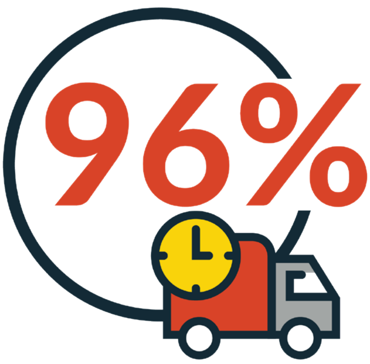

of missed time windows result in negative reviews from customers

Existing Solutions & Heuristics

Straight Line Heuristics

In an attempt to quickly solve the Vehicle Routing Problem, straight-line heuristics (or "bird paths") are commonly deployed in VRP algorithms. The use of straight-line heuristics greatly simplifies the problem by reducing the need to call external API services to determine road distances. This makes the solving process very fast and cheap. However, overreliance on straight-line heuristics may lead to unnatural routes for drivers, especially in contexts where there are physical obstructions between two distant points. For instance, if a routing engine makes use of only straight-line heuristics for stops littered on different banks of the river, the engine may produce a route that instructs the driver to cross the river back and forth instead of serving as many stops on the same side of the riverbank without crossing the same bridge repeatedly. An algorithm overly reliant on straight-line heuristics loses awareness of the topographical context.

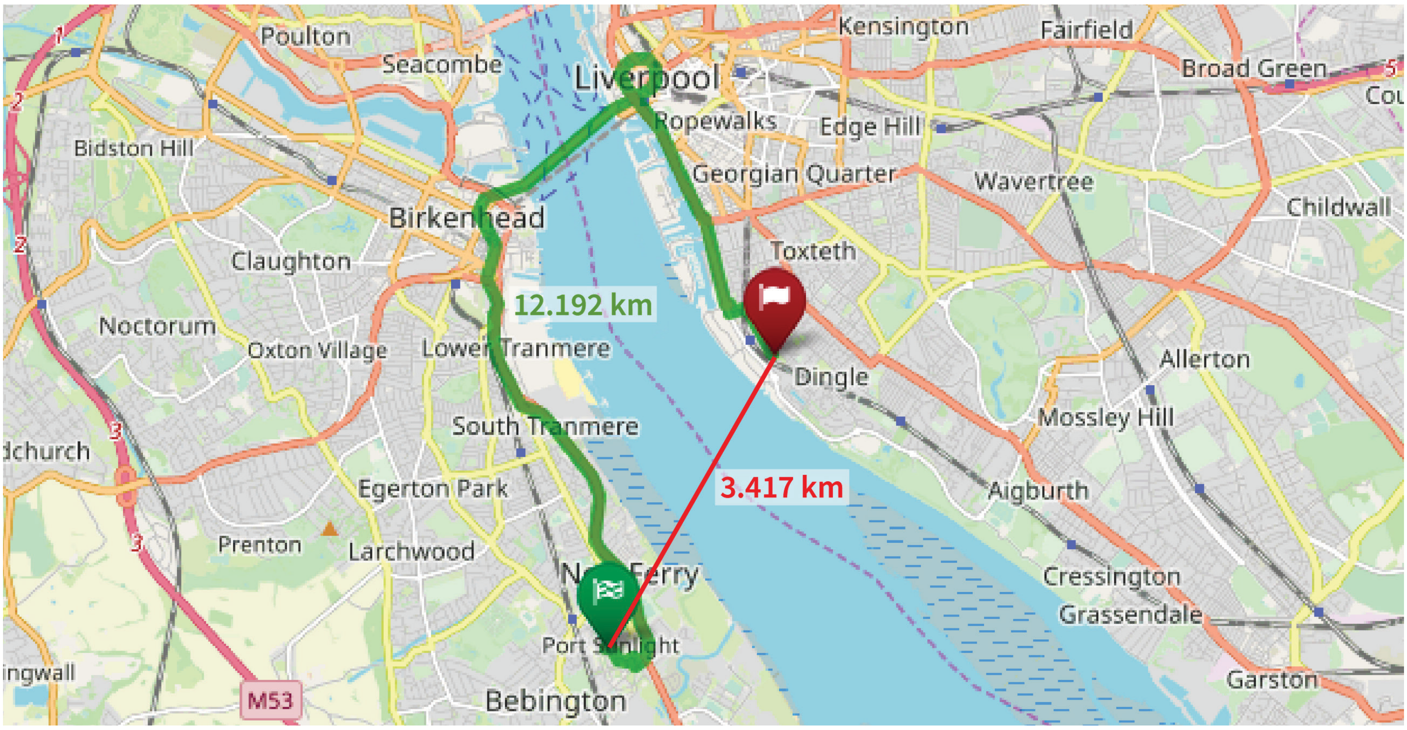

Figure 4: Real Road Distance metric versus Bird Path (straight line) heuristics. Over-reliance on straight-line heuristics may create routes that instruct drivers to cross rivers like these an impractical number of times.

Real Road Distances

The alternative solution would then be to use real road distances as a metric for every calculation, but as discussed in the TSP section, this method would be very slow and costly. Since one has to contact external services such as Google Maps for this information, API calls like these are slow and not scalable enough to solve VRP's in the magnitude of thousands of stops and above.

Number of Stops	Number of API Calls	Price (USD)
100	10000	50
500	250000	1000
1000	1000000	4100

Table 2: The cost of building the distance matrix for a given number of stops if one were to use real road distance metrics for all the stops, using the Google Maps API

It can be seen that neither straight line heuristics nor real road distance metrics alone can produce acceptable results for route planners in a reasonable amount of time in the context of Last-Mile Delivery. In the next section, we introduce a unique method that can solve large scale Vehicle Routing Problems at speeds never seen before in the industry.

The Solution: Hybrid Heuristics

We propose a solution that utilizes both the straight-line heuristics and real road distance metrics to create a routing algorithm that can solve large scale Vehicle Routing Problems with highly accurate approximations in a short amount of time, while producing natural routes and staying aware of topographical contexts such as rivers and mountains. The most significant merits of this hybrid heuristics approach are being able to produce natural clustering and sequences at a high speed.

Hybrid Heuristics

Neither pure straight-line heuristic nor a pure road-distance metric solution would suffice in solving real life Last-Mile Delivery problems. With hybrid heuristics, we use straight-line heuristics to make smart guesses about which road-distance metrics are worth evaluating, to reduce the number of times we make external API calls to evaluate road distances. This gives us the accuracy benefit of road-distance metrics without the burden of slow speeds and high costs. With this increased speed, it also gives us more room to explore different clustering techniques in parallel to see which gives the greatest savings in total journey time/distance for each vehicle in the VRP.

Natural Clustering

Manual zoning is labor-intensive and cannot react well to unexpected surges or variations in demand across different geographical zones. With a machine powered clustering algorithm, the routing engine automatically creates micro-clusters based on the stops, not the geographical region, and then assigns vehicles to serve these micro-clusters. This allows the algorithm to naturally respond to sudden surges in demand. When stops appear in a certain geographical zone in greater numbers than usual, the algorithm automatically draws out more clusters in the region, allocating more vehicles to serve this region without manual human intervention.

Natural Clustering Investigated

In this section, we will investigate how natural clustering performs against hard zoning using the same dataset of 100 stops and 4 vehicles. In this example, we divide mainland Singapore into four zones, North, South, East and West, and assign one vehicle per zone. The vehicle must serve all stops inside its zone within the time window of 9am to 5pm.

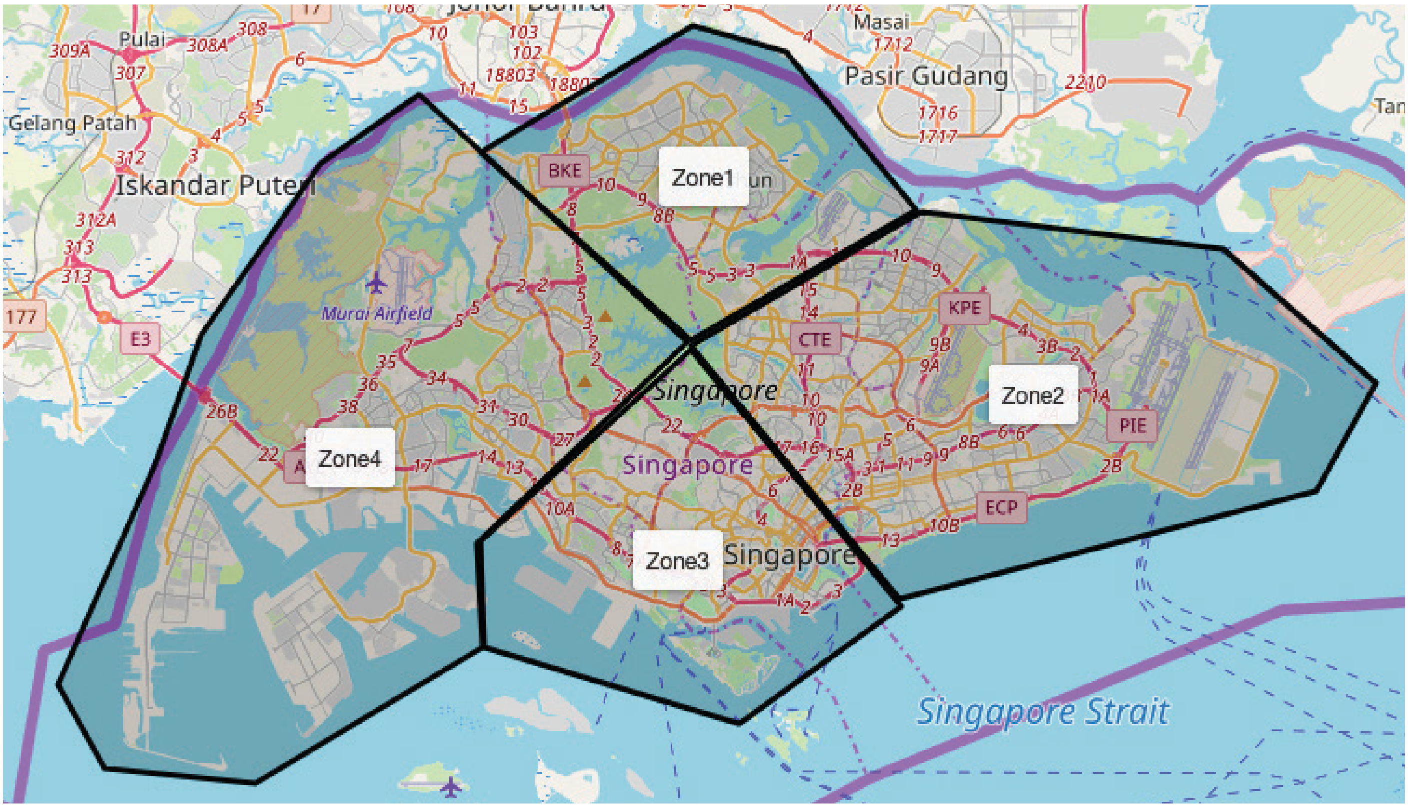

Figure 5: A map of Singapore divided into four geographical zones for the purpose of assigning vehicles to stops

We load a sample of 100 stops randomly scattered across residential areas in all four zones. In this distribution, the East Zone (Zone 2 in Figure 6) has a higher concentration of stops. We see that it is physically impossible for the lone vehicle in the East Zone to serve all the stops in the Zone.

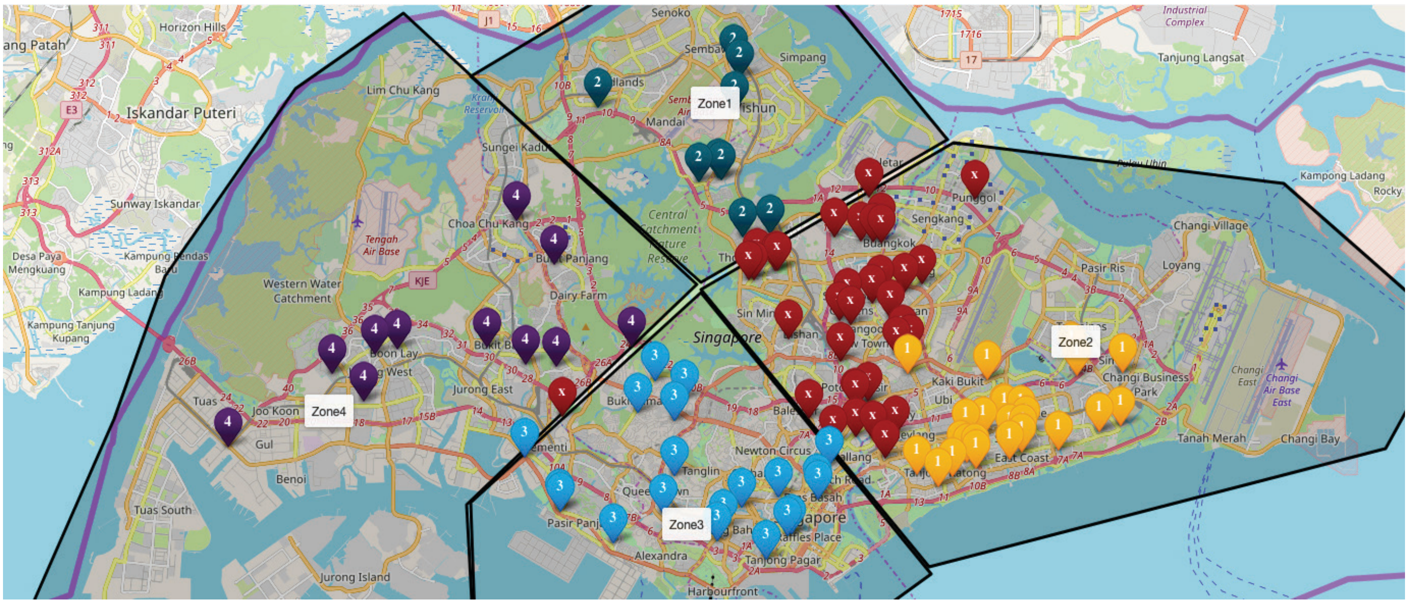

Figure 6: Results of the planning using hard zones. Stops that cannot be served are marked in red

Evidently, these hand-drawn zones are not flexible enough to respond to variances in the density of stops across different geographical areas. In Figure 6, we see that the vehicle in North Zone (Zone 1) only serves 8 stops, while the vehicle in East Zone (Zone 2) cannot serve all the stops in time. The inefficiency of the hard zoning approach to routing is illustrated by the number of unserved stops (32) within the zone. The performance of the routing algorithm is now dependent on how evenly distributed the stops are across the different zones, which is highly variable and cannot be controlled.

With natural clustering however, instead of hand drawn zones, the algorithm derives the zones out of the stops themselves. Upon removing the hard limit imposed by the zones, we can see the routing algorithm is able to evenly distribute work across all vehicles (Figure 7) without leaving any stop unserved.

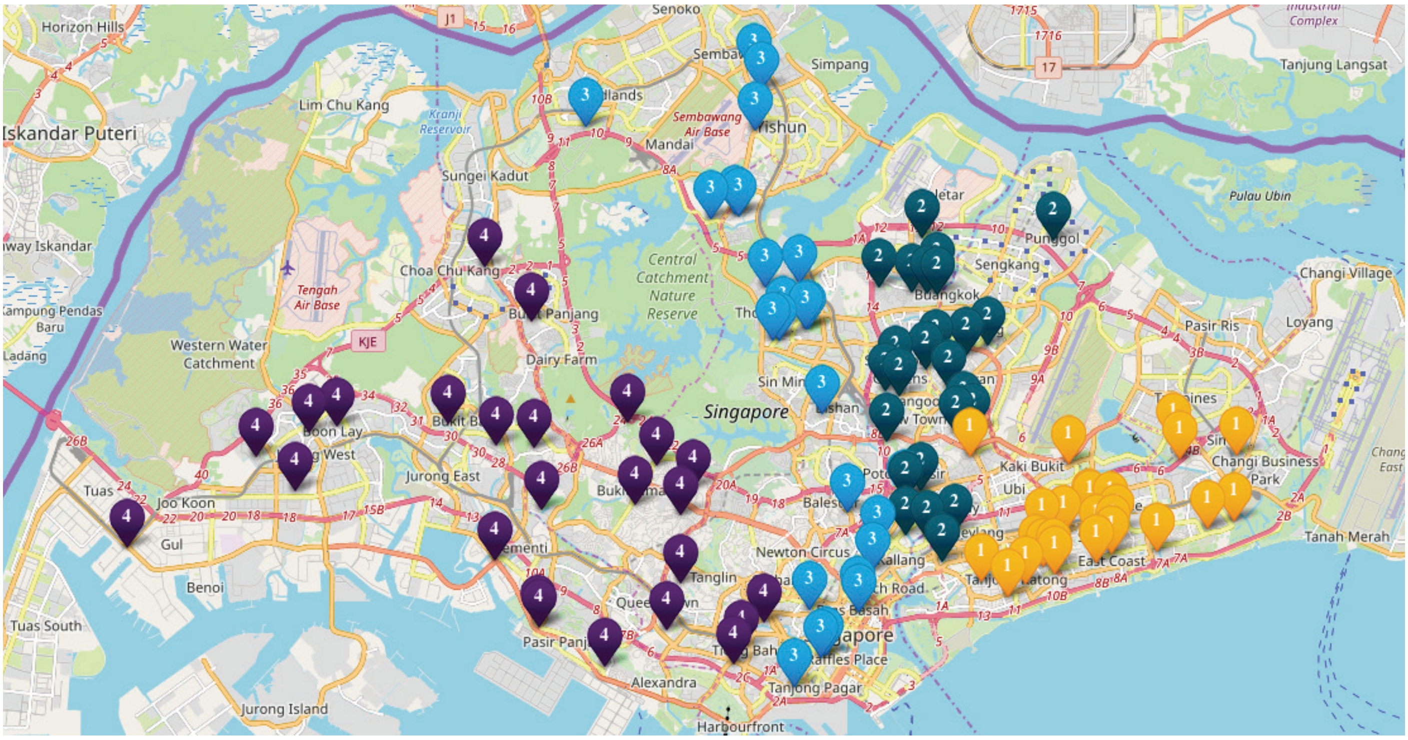

Figure 7: Natural clustering in action

However, the process of clustering the stops is very challenging and can also get computationally expensive, because there exists so many different combinations of clusters to evaluate. Hybrid Heuristics takes the place of traditional expensive metrics in the clustering process to help to reduce the computation time with smart approximations without greatly sacrificing on accuracy, making the natural clustering process much faster.

Natural Flow & Sequences

Clustering helps the routing engine cover as many stops as possible within time constraints, but the natural flow of the routes within the cluster is also needed to save even more traveling time and present intuitive routes that are easy to follow for drivers. While producing routes with a natural flow is definitely possible with traditional approaches, it is too slow and expensive to be used in a business environment. Hybrid Heuristics makes the process of finding the natural routes much faster and more practical to be used in Last-Mile Delivery. As a result, the algorithm, using Hybrid Heuristics, is able to quickly solve VRP's without having the zig-zag, unnatural routes that greedy algorithms commonly exhibit. An example is shown below with several stops in the city of Liverpool.

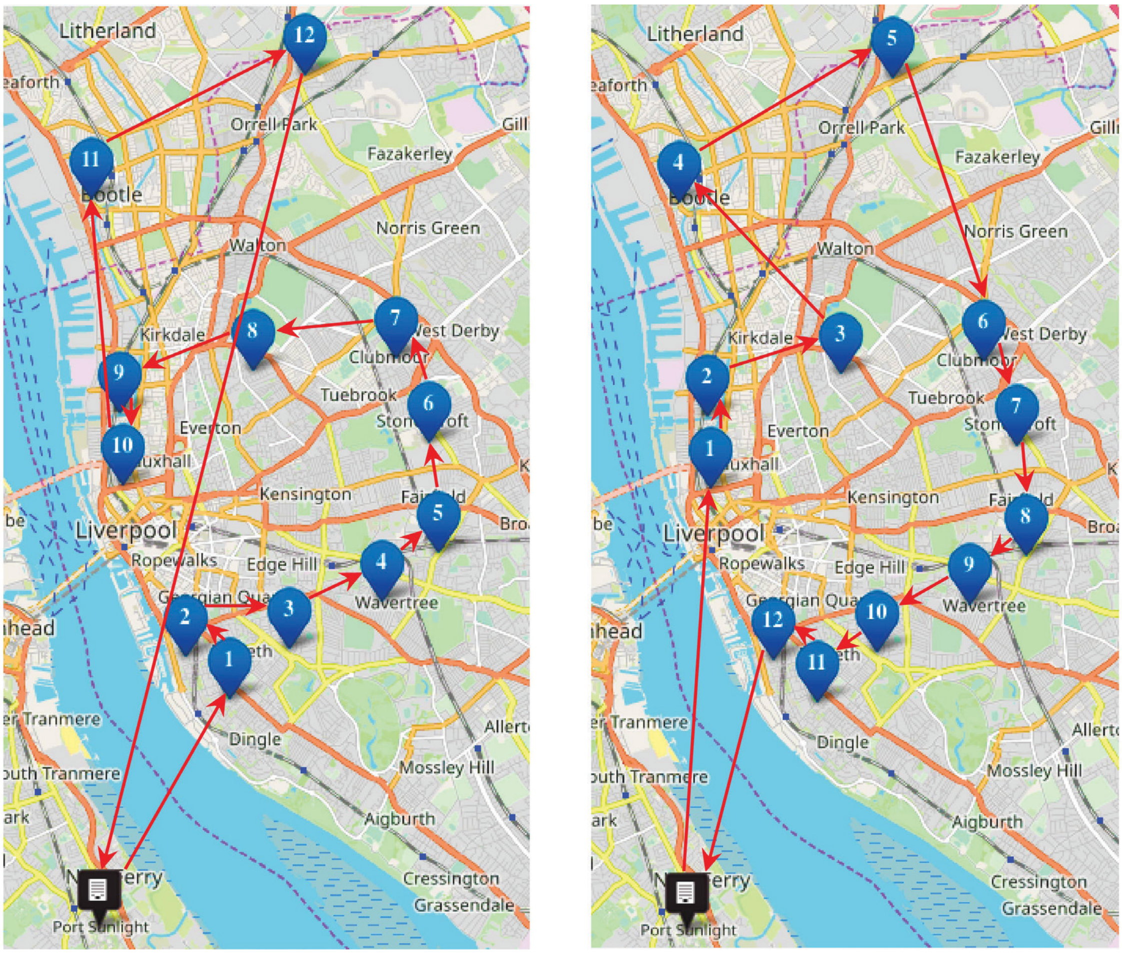

Figure 8: Sequences of a greedy algorithm (left), and of a natural flow produced at high speed using hybrid heuristics (right). The arrows in red serve as a visual guide of the direction of flow.

Natural routes that reduce zig-zag maneuvers also generally reduce the road distance travelled – for instance, the route that the greedy algorithm produces has a road distance of 58 km, while the optimized route has a road distance of 53 km (See the appendix for a table of the addresses of the stops).

Tangible Performance Benefits

Number of Stops	Number of Vehicles	Time Taken (secs)	Time Taken (mins)
100	3	17	0.28
500	10	86	1.43
1000	19	188	3.13
2000	33	422	7.03
5000	78	1474	24.56

Table 3: Performance Benchmarks of the ElasticRoute Engine

With new insights gathered from studying and refining these hybrid heuristics, we launched a new service, ElasticRoute, that can solve large scale Vehicle Routing Problems within an improved time frame that is required in the business context of Last-Mile Delivery, while providing the advantages of natural clustering and naturally flowing routes.

The use of Hybrid Heuristics has proven to give tangible performance benefits: ElasticRoute is able to solve a 1000-stop VRP under 5 minutes and a 5000-stop VRP under 25 minutes (Table 3). With a faster, cheaper and more scalable routing engine, couriers and route planners are now able to accommodate for last-minute changes to the manifest and complete the route plans way ahead of time before the drivers need to be dispatched without incurring a large cost and time from re-planning.

The earlier mentioned “choose two out of three” problem: high speed, natural routes, and low computation costs, is no longer a problem with ElasticRoute. Couriers can leverage the ElasticRoute advantage with its Hybrid Heuristics to gain a footing in the Last-Mile Delivery scene of tomorrow, today.

Appendix

Name	Address	Latitude	Longitude	Road Distance between the Stops (km)
Stop12	1 Greaves Street, Liverpool	53.38527371	-2.9693985	12.176
Stop9	1 Carmarthen Crescent, Liverpool	53.39069973	-2.97901154	1.383
Stop7	1 Jermyn Street, Liverpool	53.39167864	-2.95633078	1.932
Stop6	1 Tabley Road, Liverpool	53.39783952	-2.93618202	2.518
Stop10	1 Rathbone Road, Liverpool	53.4041658	-2.92362392	2.814
Stop4	1 Carnegie Road, Liverpool	53.41853558	-2.92579651	3.255
Stop3	1 Brookbridge Road, Liverpool	53.42968527	-2.93307214	1.539
Stop1	1 Maslin Drive, Liverpool	53.4280488	-2.964077	2.857
Stop5	1 Houlgrave Road, Liverpool	53.4224741	-2.99325943	2.897
Stop11	Leeds Street, Liverpool	53.41295324	-2.99287319	1.789
Stop2	1 Irlam Road, Bootle	53.45023805	-2.9996109	5.215
Stop8	1 Goodacre Road, Liverpool	53.46643808	-2.95291901	4.161
Return to Depot	1 Windy Bank, Wirral	53.35516211	-2.99789429	15.787
			Total (km)	58.323

Table 4: Journey Log of the Greedy Route

Name	Address	Latitude	Longitude	Road Distance between the Stops (km)
Stop11	Leeds Street, Liverpool	53.41295324	-2.99287319	9.5
Stop5	1 Houlgrave Road, Liverpool	53.4224741	-2.99325943	1.686
Stop1	1 Maslin Drive, Liverpool	53.4280488	-2.964077	2.901
Stop2	1 Irlam Road, Bootle	53.45023805	-2.9996109	4.388
Stop8	1 Goodacre Road, Liverpool	53.46643808	-2.95291901	4.161
Stop3	1 Brookbridge Road, Liverpool	53.42968527	-2.93307214	5.032
Stop4	1 Carnegie Road, Liverpool	53.41853558	-2.92579651	1.632
Stop10	1 Rathbone Road, Liverpool	53.4041658	-2.92362392	3.261
Stop6	1 Tabley Road, Liverpool	53.39783952	-2.93618202	2.901
Stop7	1 Jermyn Street, Liverpool	53.39167864	-2.95633078	2.452
Stop12	1 Greaves Street, Liverpool	53.38527371	-2.9693985	1.664
Stop9	1 Carmarthen Crescent, Liverpool	53.39069973	-2.97901154	1.383
Return to Depot	1 Windy Bank, Wirral	53.35516211	-2.99789429	12.122
			Total (km)	53.083

Table 5: Journey Log of the Optimized Route

Credits

Produced by Detrack Systems, September 2019

Author

***Chester** Koh*

Editors

***Aline** Tan*
***Charles** Tan*
***Johnny** Tan*

Graphic Designers & Typesetters

***Farah** Nur Afrina*
***Gabrielline** Oliveiro*
***Iqbal** Ghani*

Reviewers

***Dason** Goh*
***Fanny** See*

From the creators of Detrack Proof of Delivery, ElasticRoute is a new service that brings fast, affordable and scalable route planning to your fleets. Find out more at www.elasticroute.com

Connect with Us

facebook.com/elasticroute

linkedin.com/company/elasticroute

[@elasticroute]()

youtube.com/c/ElasticRoute

Media Enquiries

Iqbal Ghani, Senior Marketing Executive
Detrack Systems Pte Ltd
iqbal@detrack.com