Building a Scalable Electrification Infrastructure in Logistics

by

Muhammad Ashhad Alam

S.B. Computer Science and Engineering
MIT, 2023

Submitted to the Department of Electrical Engineering and Computer Science
in partial fulfillment of the requirements for the degree of

MASTER OF ENGINEERING IN ELECTRICAL ENGINEERING AND COMPUTER SCIENCE

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

September 2024

Authored by: Muhammad Ashhad Alam
Department of Electrical Engineering and Computer Science
June 30, 2024

Certified by: Alexandre Jacquillat
Associate Professor, Operations Research and Statistics, Thesis Supervisor

Accepted by: Katrina LaCurts
Chair
Master of Engineering Thesis Committee

Building a Scalable Electrification Infrastructure in Logistics

by

Muhammad Ashhad Alam

Submitted to the Department of Electrical Engineering and Computer Science
on June 30, 2024 in partial fulfillment of the requirements for the degree of

MASTER OF ENGINEERING IN ELECTRICAL ENGINEERING AND COMPUTER SCIENCE

ABSTRACT

The transportation sector in the US contributes to about a third of all greenhouse gas emissions, about a quarter of which stems from road freight. A major driver of this environmental footprint remains a heavy reliance on trucking—the least fuel-efficient mode of transportation. A key pathway toward freight decarbonization, therefore, involves shifting from internal combustion engines (ICE) to electric powertrains in truck fleets. This work develops analytics-based solutions to support and assess the electrification of long-haul logistics operations, by applying the methods to PepsiCo's operations in Texas.

Thesis supervisor: Alexandre Jacquillat
Title: Associate Professor, Operations Research and Statistics

Acknowledgments

I would first and foremost like to show my appreciation for Prof. Alexandre Jacquillat. This research would not have been possible without his constant help, and guidance. I especially appreciated his constant critiquing of any methods proposed, which helped prevent me from getting distracted by potentially less useful methods, as well as improve the ones I did end up using.

I would also like to thank Sean Lo for his invaluable assistance. Sean not only helped keep me on track, but was always there to exchange ideas, and help out with blockers. In addition, I will always appreciate the time he took out to help me debug the model when I would get stuck on it.

Moreover, I would like to acknowledge Andrea Zanon, who's initial ideas I built on to tackle the problems outlined in this thesis.

Lastly, I would also like to thank PepsiCo for providing the dataset used for this analysis.

Contents

Title page 1

Abstract 3

Acknowledgments 5

List of Figures 9

List of Tables 11

1 Introduction 13
- 1.1 Electric Vehicles and Climate Change . . . 13
  - 1.1.1 Comparison with Internal Combustion Engines . . . 13
- 1.2 Barriers to Freight Electrification . . . 13
  - 1.2.1 Operational Reliability . . . 14
  - 1.2.2 Managerial Barriers . . . 14
  - 1.2.3 Infrastructural Barriers . . . 14
- 1.3 Drivers for Freight Electrification . . . 14
  - 1.3.1 Public Policy . . . 15
  - 1.3.2 Cost Benefits . . . 15

2 Literature Review 17
- 2.1 Related Work . . . 17
- 2.2 Methodological Literature Review . . . 17

3 Model 19
- 3.1 Dataset . . . 19
  - 3.1.1 Assumptions . . . 20
  - 3.1.2 Pre-Processing . . . 21
  - 3.1.3 Costs of Operation . . . 24
- 3.2 Problem Statement . . . 24
- 3.3 Model Formulation . . . 25
  - 3.3.1 Parameters . . . 25
  - 3.3.2 Decision Variables . . . 26
  - 3.3.3 Stage One . . . 27
  - 3.3.4 Stage Two . . . 28

4 Model Analysis **31**
4.1 Pre-Processing . . . 31
4.1.1 Domination Criteria for Edge Combinations . . . 32
4.2 Model . . . 32

5 Results and Practical Takeaways **35**
5.1 Results . . . 35
5.1.1 Baseline Algorithm (Greedy) . . . 35
5.2 Practical Takeaways . . . 39
5.2.1 Insights . . . 41

6 Conclusion **45**
6.1 Future Work/Next Steps . . . 45

A Full Results **47**
A.1 Electrification of the Texas Triangle by the Baseline . . . 47
A.2 Comparison to Baseline . . . 48
A.3 Electrification as Budget Increases . . . 49

References **51**

List of Figures

3.1 Best-Case and Worst-Case Range for Electric Vehicles from the 4 Main Cities 20
3.2 Cities Picked During Pre-Processing 21
3.3 Delaunay Triangulation of the Fully Connected 4 City Graph 22
3.4 Example Edge Combination to Electrify 23
3.5 Number of Stations Required to Make Each Edge Feasible in Example 23
3.6 Shortest Paths between All Nodes in Example 24

4.1 Variation in Computation Time versus Number of Electic Vehicles 33
4.2 Variation in Computation Time versus Budget 33
4.3 Variation in Computation Time versus Number of Orders 34

5.1 Comparison of Model to Baseline in the 4 City Case for Budget 7 36
5.2 Comparing Model Against the Baseline for 4 Cities 37
5.3 Baseline Electrification of the Texas Triangle as Budget Increases in the 10 City Case 37
5.4 Model Electrification of the Texas Triangle as Budget Increases in the 10 City Case 38
5.5 Comparing Model Against the Baseline for 10 Cities 38
5.6 Flow of Traffic Between the Different Cities 39
5.7 Electrification of the Texas Triangle as Budget Increases in the 4 City Case . 40
5.8 Electrification as Budget Increases in the 10 City Case in a Majority EV Fleet 40
5.9 Electrification of the 10 City Case with a Budget of 5 41
5.10 Percentage of Demand Met as Budget Varies 42

A.1 Electrification of the Texas Triangle as Budget Increases in the 4 City Case . 47
A.2 Baseline Electrification of the Texas Triangle as Budget Increases in the 10 City Case 47
A.3 All Results, 4 Cities 48
A.4 All Results, 10 Cities 48
A.5 Electrification as Budget Increases in the 10 City Case in a Majority ICE Fleet 49

List of Tables

3.1 Best-Case and Worst-Case Scenarios of Electric Truck Technologies 20

Chapter 1

Introduction

Greenhouse Gases (GHGs) are one of the leading causes of climate change, and in particular, global warming. According to the United States Environmental Protection Agency, in 2022, the Transportation Sector contributed to about 28% of all GHG emissions, about a quarter of which can be attributed to medium- and heavy-duty (MHD) trucks [1].

1.1 Electric Vehicles and Climate Change

Electric vehicles (EVs) are considered a cleaner alternative to ICEs. Most vehicles are powered by an Internal Combustion Engine (ICE), which burns fossil fuels and releases GHGs into the atmosphere. EVs, on the other hand, run on batteries that power the electric motors, which is better for the environment.

1.1.1 Comparison with Internal Combustion Engines

While EVs are better for the environment than ICEs, they still generate some pollution. The process of extracting and processing the minerals required for the batteries, as well as manufacturing the batteries, generates a lot more carbon footprint than the manufacturing of ICEs [2].

In addition, depending on how the electricity to power them is generated, the operation of EVs also generates some pollution. However, even if the electricity was generated from fossil fuel-based power stations, EVs generate less pollution throughout their lives than ICEs that have been driven the same amount. In Massachusetts, based on the current grid mix (which is about 80% fossil fuel based on data from the US Energy Information Administration), electric trucks could reach their "break-even" point in less than two years [3]. As the grid keeps moving towards greener alternatives, this break-even point could arrive even sooner, and the payoffs of shifting to an electric fleet for the environment will be even greater.

1.2 Barriers to Freight Electrification

While vehicle electrification is underway in car markets, it has been lagging in freight due to several barriers [4], [5].

1.2.1 Operational Reliability

Electric powertrains are more challenging to embed into MHD vehicles due to range and payload requirements. The range and operational speed of the trucks are dependent on many factors, including weather and total weight being transported. This might lead to added costs that appear in the form of delayed deliveries and extra charging costs just to name a few issues.

As a result, the trucks would need to be reloaded more often, or we would need a larger fleet size to accommodate existing markets. In addition, the range of electric trucks is more limited and uncertain when compared to ICEs, and overestimating the range might result in added costs of towing them to the nearest charging stations.

In addition, the batteries of the electric trucks weigh a lot more than the gas tanks. However, this is not as big a problem as it seems unless the cargo is very dense. While the added weight of the battery reduces the weight of the cargo that a truck can carry [6], this is not a concern in practice. Trucks tend to "volume out" before they "weigh out" [7]. In other words, freight trucks frequently reach their volume limit for the cargo before reaching the weight limit. So, a lower weight limit is not necessarily a problem for freight trucks.

1.2.2 Managerial Barriers

Existing routes used by ICE trucks might not be the most optimal for electric trucks. In addition, we have a lot more constraints to keep track of now, such as battery life, making sure we have enough range to get to the next charging station, etc. This makes the routing problem a lot harder, and using old methods might not be the most effective way to develop new distribution networks. This is where the research in this thesis comes in; to develop a new analytics-based approach to solve this problem.

1.2.3 Infrastructural Barriers

The lack of charging infrastructure along long-haul routes is another big barrier to switching to electric trucks. Moreover, charging infrastructure features limited interoperability, which can hinder adoption in a market featuring dozens of electric truck models with specific charging requirements. To make things worse, these challenges reinforce each other, creating “chicken and egg” dynamics between electric truck manufacturing, electric truck adoption, and charging station deployment.

Furthermore, even if an "optimal placement of charging stations" were to be found, it might not be practical to build the stations at those locations due to certain factors. The geographical landscape, costs of the land itself as well as obtaining permits are just a few of these factors that can hinder this development. These infrastructural barriers significantly restrict the candidate charging stations that can be built.

1.3 Drivers for Freight Electrification

Some factors do encourage investment into freight electrification [5].

1.3.1 Public Policy

To encourage lower emissions, 15 states in the US, led by California in 2020, have already adopted, or are in the process of adopting the Advanced Clean Trucks (ACT) Regulation. The ACT Regulation requires manufacturers of MHD vehicles to increase the percentage of sales of zero-emission vehicles and near-zero emission vehicles, aiming to achieve 100% electric truck sales by 2050.

1.3.2 Cost Benefits

Saving Costs for Large Logistics Providers

Large logistics providers can play a leading role in truck electrification. While the cost of charging infrastructure can be expensive for individual drivers, large-scale logistics providers may be able to recoup the investments more easily due to higher utilization. At the same time, these opportunities create new challenges: how can a logistics provider effectively utilize a partially electrified fleet; how can a logistics provider deploy charging stations to power its fleet while saving on costs?

Saving Environmental Costs

As mentioned in Section 1.1, electric vehicles tend to be worse for the environment in the short-term as compared to ICEs [2]. However, they have significant long-term environmental benefits.

There are both long-term capital and environmental benefits to be gained by transitioning to electric vehicles. This leads to the following question: can identify enough medium-term gains to recoup capital expenses with cost savings and reductions in GHG emissions to encourage logistics decarbonization?

This thesis aims to help answer that question by developing an infrastructure that maximizes the benefits gained while minimizing the costs incurred during this transition to logistics providers.

Chapter 2

Literature Review

2.1 Related Work

Past work in this domain has primarily focused on exploring whether the electrification of freight vehicles is viable in the first place. In particular, a couple of the papers mentioned earlier discuss the different factors that influence electrification [4], [5]. The authors mainly found that most benefits from electrification (both in terms of operating costs and environmental) are long-term, while the barriers to entry are more short-term and immediate, the biggest being the current lack of infrastructural support for electric freight vehicles in most of the US.

Other relevant research in this domain includes Nicolaides et al., 2018 [8] which talks about charge-on-the-move infrastructure. This focuses on vehicles with smaller batteries and tries to provide power to them while they're still moving. While this might reduce some infrastructural costs as we might not need to build charging stations for them, currently such vehicles cannot carry enough loads to be viable options for MHD vehicles. If, however, there is more development done with charge-on-the-move vehicles, they could very well become an alternative option.

Jaller et al., 2020 [9], on the other hand, primarily focus on the different challenges we might face, and the different opportunities, such as policies that can be developed, to incentivize the electrification of the fleets further. Specifically, the authors talk about employing different technologies to aid the move to a more electric world, both for personal, as well as large-scale transport use.

2.2 Methodological Literature Review

Moreover, there has been a lot of work done that is relevant to the methodology used in this thesis. Farahani et al., 2013 [10] present strategies to tackle some network design problems specifically when it comes to transportation. Ukkusuri et al., 2008 [11] formulate a Dynamic User Equilibrium Network Design Problem using a Linear Programming approach. Both of these papers provide the basic ideas that we use to develop the graph that our linear programming model works on.

In addition, there has been some work done to tackle both routing problems, as well as routing specifically for electric vehicles. Ukkusuri et al., 2007 [12] is one such paper that addresses the problem of traffic routing under demand uncertainty. Mrazovic et al., 2018 [13] develop a multi-vehicle route planner to optimize urban freight transport. Morlock et al., 2019 [14], on the other hand, tackle the problem of specifically routing electric vehicles while keeping the charging infrastructure in mind. We build on the ideas mentioned in these papers to support the routing of electric freight trucks.

This project develops analytical solutions to support and assess the electrification of long-haul logistics operations. Research has been underway to develop optimization methods that support logistics operations with electric vehicles. This project opens a second area of research to integrate logistics operations into the design of a network of charging infrastructure. This research is conducted in close partnership with PepsiCo, to evaluate our methods on real-world data, to derive insights collaboratively on the potential for logistics electrification in PepsiCo's operations in Texas, and to work together toward a pilot deployment in practice.

Chapter 3

Model

The electrification of long-haul logistics is complicated by interdependencies between strategic infrastructure deployment and tactical logistics operations: logistics routes depend on available charging stations and, vice versa, investments in charging infrastructure depend on how they can be utilized by logistics providers. This project develops an integrated approach to jointly optimize these two sets of decisions by combining vehicle routing algorithms to support electrified operations and facility location models to guide investments in charging infrastructure.

3.1 Dataset

The dataset used for developing the model and analyzing the results in this thesis comes from PepsiCo's data, which outlines their current logistics operations. This thesis focuses on the operations in the Texas region and specifically concentrates on the area around Texas' four main cities; Austin, Dallas, Houston, and San Antonio, also known as the Texas Triangle.

The data contains a lot of information about the information relevant to the different Origin-Destination pairs (O-D pairs) for each of the orders, including:

- **Origin City**: The city where the order starts at. For larger cities, the data contains more specific information about the location of the facility the order originates from.
- **Destination City**: The city where the order is being delivered to. As with the origin cities, for larger cities, the data contains more specific information about the location of the destination facility.
- **Estimated Annual Volume**: An estimate of how much load was transported between this specific O-D pair over the entire year.
- **Number of Weeks of Volume**: The total number of weeks to fulfill this order.

The dataset is then pre-processed to generate estimated weekly loads for the O-D pairs and then reduce the data to only the cities in the Texas Triangle region.

In addition, we have gathered information on trucking technologies from previous work [15] and publicly available sources. This provides a view of the best-case and worst-case

Table 3.1: Best-Case and Worst-Case Scenarios of Electric Truck Technologies

Parameter	Description	Min	Max
MPGe	Average Fuel Economy, Electric	12.6	17.7
Charging Power	Measured in kW, determines charging speed	350	350
Battery Capacity	Commonly measured in kWh	250	450

scenarios regarding electric truck technologies (see Table 3.1). The values in the table are sourced from the National Renewable Energy Laboratory [16], and the ranges among several class 8 long-haul trucks [17].

To develop a more meaningful perspective for logistics operations, we generate the visualizations (see Figures 3.1a-3.1d) below to indicate estimates of the worst-case and best-case range of electric trucks (corresponding to battery capacities of 250 kWh and 450 kWh, respectively), starting from each of the four main cities in Texas. These patterns indicate that the range may be insufficient to cover all operations in the state, and so, highlight the importance of investing in charging infrastructure along long-haul routes.

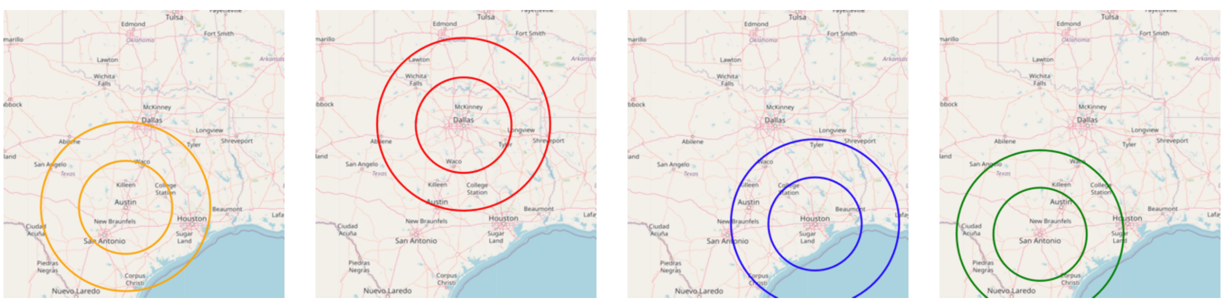

(a) Austin (b) Dallas (c) Houston (d) San Antonio

Figure 3.1: Best-Case and Worst-Case Range for Electric Vehicles from the 4 Main Cities

3.1.1 Assumptions

We make the following list of assumptions in this model:

- Each facility location has a charging station. This is a practical assumption, as logistics providers need to have warehouses in cities. Having charging stations here comes with the added benefit of saving time, as trucks can be charged while being loaded or unloaded. In addition, any trucks passing through the city can also use these charging stations.

  This further helps us assume that any edge is feasible as long as we can start at one end of the edge fully charged and traverse the edge without ever going below a battery level of 0%.

- The electric vehicles are not affected by any external factors, such as weather and cargo weight. This helps us assume a constant depletion rate of the battery, regardless of these factors.

- All trucks (ICE and EV) travel at a constant speed.

3.1.2 Pre-Processing

After reducing the dataset to only the cities in the Texas Triangle region, two samples of size 4 and 10 (see Figures 3.2a and 3.2b, respectively) are picked for analysis. The metric used to pick these cities is the total number of orders the cities are involved in, either as the origin, or the destination city.

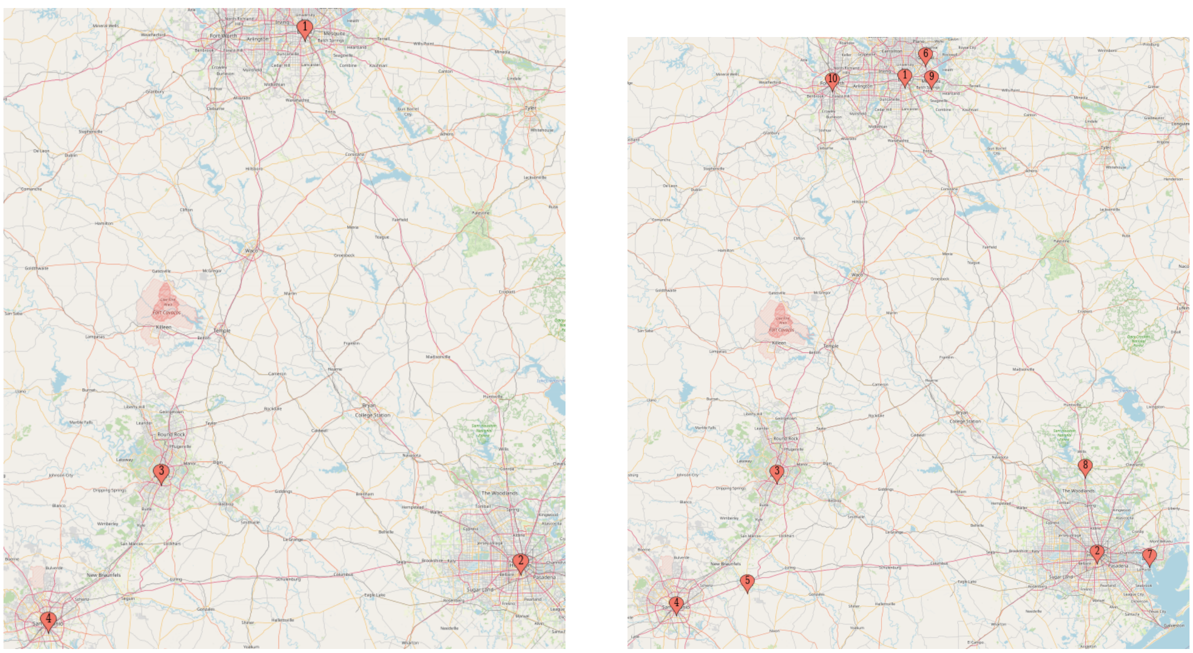

(a) 4 Cities

(b) 10 Cities

Figure 3.2: Cities Picked During Pre-Processing

Once the cities have been picked, we generate the graph the model will work on.

Generating the Graph

The graph generation involves two main steps:

1. Generate the fully connected graph using the cities picked.

2. Delaunay Triangulation to reduce the total number of edges in the graph.

Generating the fully connected graph: The vertices in the graph represent the cities chosen, and the edges represent connections between them. The edge weights represent the distance, in miles, for the edge in question. These edge weights are calculated by estimating the distance between the cities using the location data. Figure 3.3a shows the fully connected graph for the 4 city case. Note that the weights mentioned are the estimated distance in miles to traverse the edge.

Delaunay Traingulation: We first use this information to create a fully connected graph for all of our cities, and then use Delaunay Triangulation [18] to pick the set of edges that will be part of the graph, G. This helps us narrow down the number of edges in the graph while preserving the closest pairs [19] (that is, if v_i and v_j are the closest pair of vertices, they will share an edge). Using this, we convert our graph from Figure 3.3a to Figure 3.3b. As mentioned earlier, the edge weights represent the distance in miles to traverse the edge.

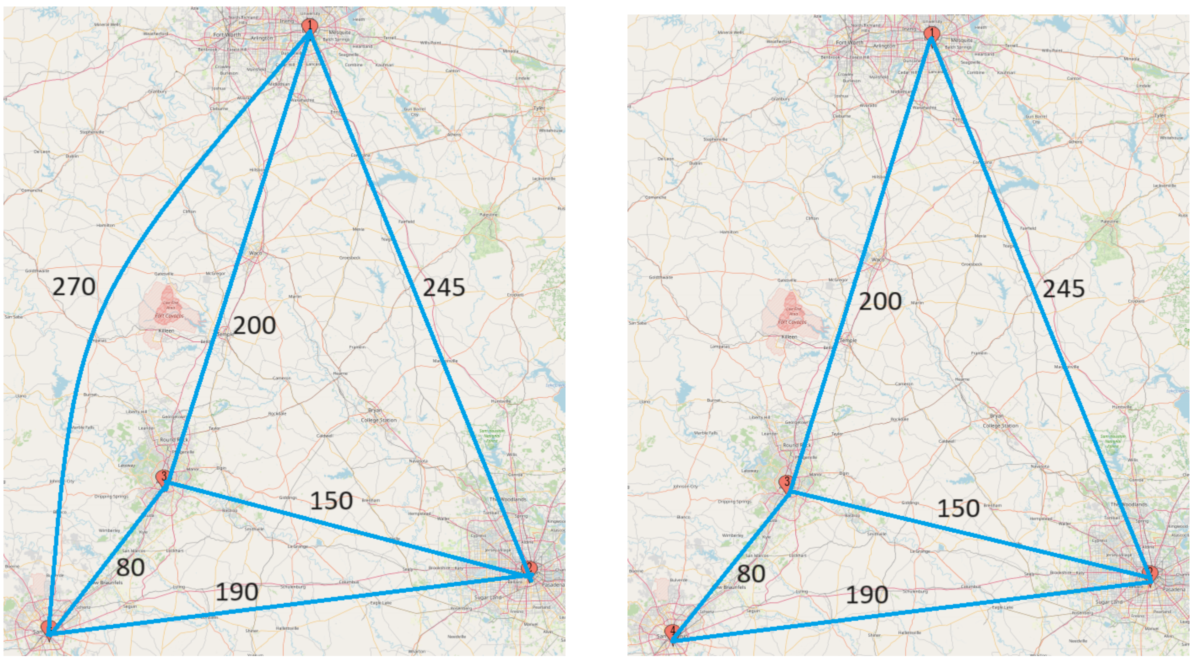

(a) Fully Connected Graph

(b) Post-Triangulation Graph

Figure 3.3: Delaunay Triangulation of the Fully Connected 4 City Graph

Pre-Processing the Possible Edge Combinations

The model wishes to determine where to optimally place charging stations on our network. Finding the precise location for each charging station is a hard enough problem to solve already, and given the infrastructure barriers (see Section 1.2.3), we might not even be able to build our candidate charging stations in those locations.

As a result, we simplify this decision problem to the much simpler problem of deciding which edges to electrify instead. Based on this choice, the model will figure out the best possible paths between all pairs of cities. To reduce computation time for the model itself, we will pre-process all possible combinations of edges that can be chosen and any other information the model requires from them. Let's look at this process using the example edge combination shown in Figure 3.4.

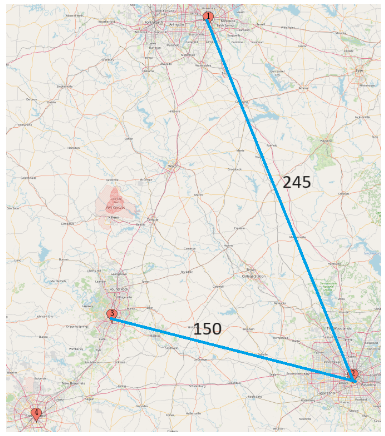

Figure 3.4: Example Edge Combination to Electrify

For the example, we will assume we are using an EV with a battery capacity of 300 kWh, and a depletion rate of 2.5 kWh/mile. This truck would have a range of about 120 miles on a full charge. On the edge combination shown, we do the following:

1. For each edge in the combination, calculate the minimum number of charging stations required to fully electrify the edge. This is calculated by using the formula:

$$num\_stations = \left\lfloor \frac{distance}{range} \right\rfloor \tag{3.1}$$

For the example, the number of stations for each edge is shown in Figure 3.5.

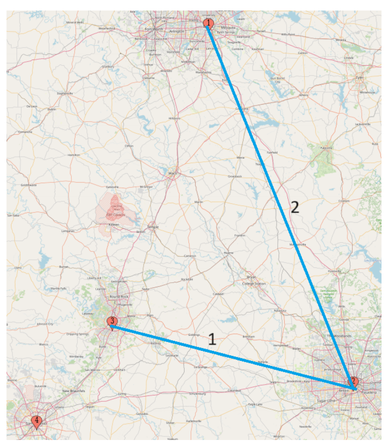

Figure 3.5: Number of Stations Required to Make Each Edge Feasible in Example

2. Sum over the costs of all individuals to find the total cost of the entire combination. In the example mentioned, we get $num\_stations = 1 + 2 = 3$ stations to make the given edge combination feasible.

3. We then run Johnson's Algorithm for All-Pairs Shortest Paths for our given edge combination. This allows the model to easily find the shortest path to fulfill any given

order using the source and destination nodes. In our example, we get the shortest paths between the 3 nodes that can be connected shown in Figure 3.6.

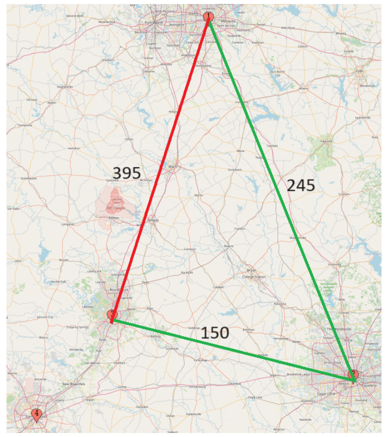

Figure 3.6: Shortest Paths between All Nodes in Example

Using the results of Johnson's we can figure out which nodes are connected, and hence, which orders our given edge combination can serve. So, for instance, in the example above, any orders to city 4 (San Antonio) cannot be transported by electric vehicles if this edge combination is chosen.

3.1.3 Costs of Operation

The costs of operation are normalized to the costs of operating an ICE vehicle. The relative cost of operation of electric vehicles is found by first calculating the cost of operation of ICE vehicles and electric vehicles and finding their ratios.

We use the average gas prices in Texas from the American Automobile Association [20] and the average mileage of freight trucks in Texas [21] to calculate the cost per mile of operating an ICE truck (0.604/mile).

Similarly, we use the ranges and battery capacities of the class 8 long-haul trucks [17], and the average electricity prices in Texas [22] to calculate the cost per mile of operating an electric vehicle (0.283-0.394/mile).

Using these two results, depending on the kind of electric truck used, we found that electric trucks on average cost 45%-65% to operate.

3.2 Problem Statement

The problem that we tackle in this thesis can be formulated as a decision problem; given our constraints (which will be discussed in more detail in Section 3.3), which routes do we electrify to maximize the number of orders fulfilled, while also minimizing the costs incurred to meet these orders.

3.3 Model Formulation

We need to allow our vehicles (trucks) to be able to serve more than one order each. Doing so requires some load-balancing as to what orders each truck should serve. To allow this without resorting to a formulation with time-indexed arcs (which suffers from the curse of dimensionality), we treat the itinerary planning of vehicles as a packing problem; we would like to assign orders to vehicles, subject to the condition that each vehicle can complete all its orders within the time horizon.

In addition, the model also needs to be able to support all kinds of mixed fleets, from fully non-electric to a hybrid (both EVs and ICEs) to a fully electric fleet. For a company deciding to electrify its fleet, an intermediate transition to a mixed fleet is a more realistic and practical option than immediately fully electrifying its fleet.

In any mixed or fully electrified fleet, we want to make sure that the operations do not suffer drastically as a result of the transition. So, we prioritize serving the demand, and then view costs of operations as a secondary objective. If any user of this model has different priorities, we can easily accommodate them by simply changing the objective function of the model.

Given our priorities, the model defined here is a two-stage optimization model for the approximate packing problem, formulated as a Mixed Integer Program (MIP). The first stage of the model outputs the maximum demand that can be met subject to the given constraints, while the second stage of the model tries to minimize the total costs incurred while trying to meet the maximum demand found in the first stage.

3.3.1 Parameters

There are quite a few parameters that will affect the performance of our model, in terms of the amount of demand met, the cost incurred, and the computation time taken to optimize the model. :

- L: set of all cities.
- I: Set of all orders, indexed by i. Each order has an origin and destination, as well as the amount of demand that can be met.
- D_i: Shorthand for maximum demand that can be met for order i.
- E: Set of all combinations in which the edges of the graph can be electrified, indexed by e.
- n_e: Number of charging stations required to electrify edge combination e.
- B: Budget, the maximum number of charging stations that can be built.
- V_{ICE}: Set of ICE vehicles.
- V_{EV}: Set of electric vehicles.
- $Days$: Set of days.

- avg_speed: The average speed of the trucks. We divide all edge distances by this parameter to get the travel time for the edges.

- t_i^{ICE}: Time required to meet unit demand for order i using ICE vehicles.

- $t_{e,i}^{EV}$: Time required to meet unit demand for order i using edge combination e using electric vehicles.

- c_i^{ICE}: Cost to meet unit demand for order i using ICE vehicles.

- $c_{e,i}^{EV}$: Cost to meet unit demand for order i using edge combination e using electric vehicles.

- T_{ICE}: Time horizon per day for ICE vehicles.

- T_{EV}: Time horizon per day for EV vehicles. Assuming that all cars start with a full charge, this can be calculated from T_{ICE} using the following formula:

$$T_{EV} = (T_{ICE} + T_{charge}) \cdot \frac{T_{travel}}{T_{charge} + T_{travel}}, \tag{3.2}$$

  where T_{charge} is the time taken to charge the battery to full, and T_{travel} is the maximum time an electric vehicle can travel on a full charge.

In addition, we also define the big M value to restrict EVs from using edge combinations that are not chosen. The total amount to be served across all orders is a sufficiently large value for this since serving any more than that gives us no added benefit and just increases our costs. As a result,

$$M = \sum_{i \in I} D_i$$

3.3.2 Decision Variables

Our model needs to make some decisions; which charging stations to electrify, how much of the demand is met using electric vehicles vs ICE vehicles, etc. As a result, our optimization model has the following decision variables:

- $EV_{v,d,e,i}$: Variable denoting the amount electric vehicle v transports on day d using edge combination e for order i.

- $ICE_{v,d,i}$: Variable denoting the amount an ICE vehicle v transports on day d for order i.

- y_e: Binary variable denoting whether edge combination e was chosen. The edge combination chosen tells us which of the edges in our graph were electrified.

- z_i: Variable denoting the demand met for order i.

3.3.3 Stage One

Objective

The first stage aims to maximize the amount of demand that can be met given our parameters and constraints. So, the objective can be written out as follows:

$$\max \quad \sum_{i \in I} z_i \tag{3.3}$$

Constraints

The first stage of the model has the following constraints:

- Maximum Demand for order i:

  $$z_i \leq D_i, \quad \forall\, i \in I \tag{3.4}$$

  This constraint provides an upper bound, D_i, for the amount, z_i, that can be delivered for every order, $i \in I$.

- Pick edge combination:

  $$\sum_{e \in E} y_e \leq 1 \tag{3.5}$$

  Since the y_es are binary variables, this constraint makes sure that we do not pick any more than 1 of the possible edge combinations.

- Budget Constraint:

  $$\sum_{e \in E} y_e \cdot n_e \leq B \tag{3.6}$$

  This constraint makes sure that the number of charging stations, n_e, required to electrify our chosen edge combination, e, does not exceed the budget, B. If y_e is 0, the combination e contributes 0 to the amount spent or charging stations, otherwise, it contributes n_e.

- Demand Served for order, i:

  $$\sum_{d \in Days} \left(\sum_{v \in V_{ICE}} ICE_{v,d,i} + \sum_{v \in V_{EV}} \sum_{e \in E} EV_{v,d,e,i} \right) \geq z_i, \quad \forall\, i \in I \tag{3.7}$$

  This constraint ensures that our decision variables for ICE vehicles, $ICE_{v,d,i}$, and EVs, $EV_{v,d,e,i}$ are meeting the demand, z_i, for order i that our model claims to meet. To do this, for each order, i, we take the sum of the amount transported over all ICE vehicles and electric vehicles over all days and make sure that it is at least as much as z_i.

- Daily Time Horizon Constraint for ICE vehicles:

$$\sum_{i \in I} ICE_{v,d,i} \cdot t_i^{ICE} \leq T_{ICE}, \quad \forall\, v \in V_{ICE},\ \forall\, d \in Days \tag{3.8}$$

  This constraint makes sure that for every ICE vehicle, v, the total time traveled on any given day, d, does not exceed T_{ICE}. The time taken for each order is the product of the time taken for each order, i, and the number of trips vehicle v makes for order i on day d. The total time spent by vehicle v on day d, therefore, is just the sum of the time spent across all the orders.

- Daily Time Horizon Constraint for electric vehicles:

$$\sum_{i \in I} \sum_{e \in E} EV_{v,d,e,i} \cdot t_{e,i}^{EV} \leq T_{EV}, \quad \forall\, v \in V_{EV},\ \forall\, d \in Days \tag{3.9}$$

  This constraint makes sure that for every EV vehicle, v, the total time traveled on any given day, d, does not exceed T_{EV}. This constraint is defined similarly to the one for ICEs, except we also take the sum over the different possible edge combinations.

- Big M Constraint:

$$\sum_{i \in I} \sum_{v \in V_{EV}} \sum_{d \in Days} EV_{v,d,e,i} \leq M \cdot y_e, \quad \forall\, e \in E \tag{3.10}$$

  This constraint ensures that electric vehicles can transport the maximum number of orders over edge combination e when $y_e = 1$, and do not transport anything over it when $y_e = 0$.

Using this set of constraints and the objective function, we get the maximum possible demand that can be met, Z, which we pass on to stage two of the model.

3.3.4 Stage Two

The stage two model is very similar to the model in stage one. As a result, this section only details the differences between the model from stage one to stage two.

Objective

Since the model is now trying to minimize the total cost to meet demand, Z, we change the objective function to be the following:

$$\min \sum_{i \in I} \sum_{d \in Days} \left(c_i^{ICE} \cdot \sum_{v \in V_{ICE}} ICE_{v,d,i} + \sum_{e \in E} c_{e,i}^{EV} \cdot \sum_{v \in V_{EV}} EV_{v,d,e,i} \right) \tag{3.11}$$

The objective function here calculates the total cost. For each order, i, we take the sum of the $cost\_per\_trip \cdot number\_of\_trips$ for both ICE and electric vehicles to find the cost incurred for this particular order. Then, we take the sum over all orders to find the total cost incurred, which the model will attempt to minimize.

Constraints

The constraints are largely the same as the ones defined in Section 3.3.3. The only change we make is to add the following constraint to the model to make sure that the model meets the maximum demand that was calculated in the first stage:

$$\sum_{i \in I} z_i = Z \tag{3.12}$$

Using this two-stage model, we will get the maximum amount of demand that can be fulfilled while minimizing the total costs incurred.

Chapter 4

Model Analysis

4.1 Pre-Processing

The complexity of each of the steps in the pre-processing is as follows:

1. **Graph Generation:** Calculating the distance between any given pairs of cities is constant time, and there are $O(|L|^2)$ such pairs, where L is the set of all locations.

2. **Delaunay Triangulation:** This part of the pre-processing happens in $O(|L| \log |L|)$ time [23].

   Since Delaunay Triangulation results in $O(|L|)$ edges [19], the resulting graph of $O(|L|)$ edges and vertices can be stored in $O(|L|)$ space.

3. **Processing the edge combinations:** This is the most time-consuming part of the pre-processing. Delaunay Triangulation has $O(|L|)$ edges, so we have a total of $O(2^{|L|})$ different edge combinations. Note that calculating (and storing) the number of charging stations required to electrify an edge takes $O(1)$ time and space, using Equation 3.1.

   Running Johnson's Algorithm on each edge combination takes $O(|L|^2 \log |L|)$ time since our graph is sparse. So, the pre-processing of edge combinations takes $O(2^{|L|} \cdot |L|^2 \log |L|)$ total time.

   For each edge combination, we store the $O(|L|^2)$ shortest paths between the pairs of vertices, and the total number of stations, which takes constant space. So, the total space complexity of this step is $O(2^{|L|} \cdot |L|^2)$ space.

The complexity of the pre-processing is dominated by the processing of the different edge combinations. So, the total time complexity of the pre-processing ends up being $O(2^{|L|} \cdot |L|^2 \log |L|)$, and the space complexity is $O(2^{|L|} \cdot |L|^2)$.

This is the worst-case complexity of the pre-processing step. Since some edges might be too short to require charging stations (especially as electric vehicles become more efficient and their range increases), we can reduce the computational time significantly by implementing domination criteria between edge combinations.

4.1.1 Domination Criteria for Edge Combinations

We say that edge combination e_1 dominates combination e_2 if the following conditions are met:

- e_1 requires at most as many charging stations to be fully electrified as e_2.
- e_2 electrifies a subset of the edges of e_1.

In such a case, e_1 is always a better option than e_2, so we can simply ignore e_2 for any of our computations.

4.2 Model

Since we have both decision variables, as well as constraints, in our model that are dependent on the number of edge combinations, in addition to the model being a Mixed Integer Linear Program (MILP), which is known to be NP-hard, the model inherently has a worst-case exponential runtime to find the most optimal solution. However, using different acceleration techniques, such as Branch-and-Bound, and Cutting Planes, we can bring down the computation time to practical levels. For the 4 city case, since the input size is small enough, we can both construct and optimize the model in under a minute in all cases.

For the 10 city case, on the other hand, the model would sometimes take longer depending on the parameters provided to it, and convergence to the optimal is not always guaranteed. As a result, we run the model with a 5% optimality gap to provide "good enough" solutions. In this section, we look at how some of the different parameters affect the computation time of the model for the 10 city case. The times stated below are the total times (in seconds) for the creation of the model and the optimization time for both stages of the model.

Number of Vehicles and Composition of Fleet: We noticed that as the number of vehicles in the model increases, the computational time rises, as expected. While it might seem like the fleet size directly impacts the computation, a further analysis of the computation time versus just the number of electric vehicles (in different compositional settings), as in Figure 4.1 shows that it is, in fact, the number of electric vehicles that affect the total computational time spent. We can see that fleets of different sizes but similar numbers of EVs tend to have similar runtimes.

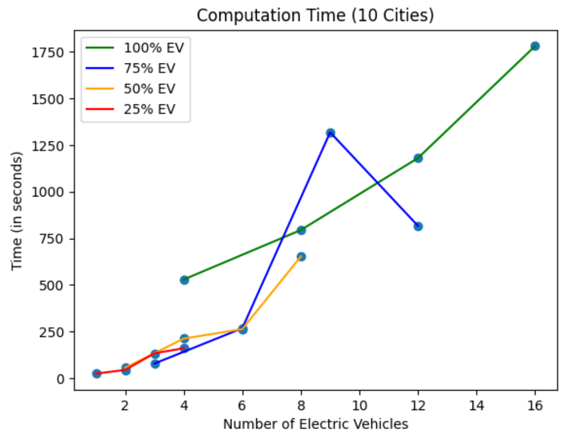

Figure 4.1: Variation in Computation Time versus Number of Electic Vehicles

Budget: The computation time is low for the lower budget ranges and also decreases closer to the higher budget range (that is when we have enough budget to electrify all edges) as demonstrated in Figure 4.2. Sometimes, the model did take a long time to converge in the intermediate budget range, which is something that did not happen in the edge range of the boundary.

At lower budget ranges, the model can eliminate several edge combinations since they are too expensive to be an option, which helps solve the problem faster. At the higher budget range, the model can pick the fully electrified (or close to fully electrified) network as the best edge combination since every other combination is a subset of it and hence, more restrictive. However, since the model has more choices to evaluate before deciding the "best" edge combination, this is still a harder problem to solve than at the lower budget values.

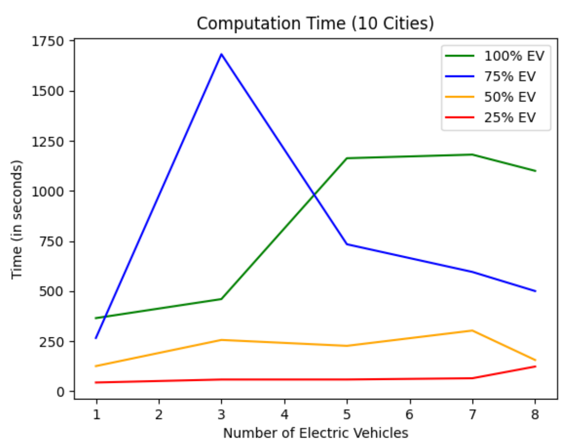

Figure 4.2: Variation in Computation Time versus Budget

Number of Orders: We notice that as the number of orders increases, the computation time of the model also increases, as shown in Figure 4.3. This can be explained by the fact that as the number of orders increases, the model has to make more routing decisions to meet the overall demand, which makes it a harder problem to solve.

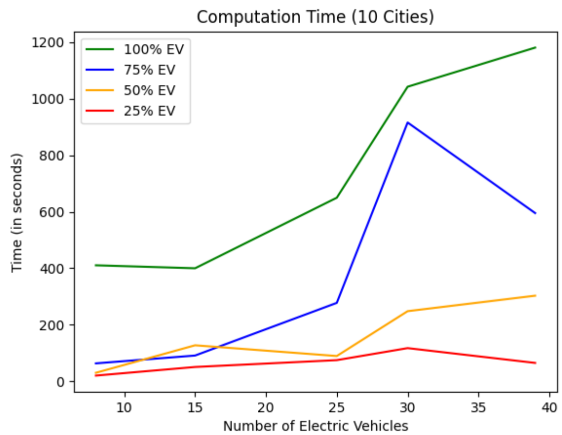

Figure 4.3: Variation in Computation Time versus Number of Orders

Chapter 5

Results and Practical Takeaways

This section of the thesis deals with our qualitative findings from the research. For each of the following sections, only the results relevant to the findings are shown. A full set of results can be found in Appendix A.

5.1 Results

The model was mainly evaluated by comparing it with a baseline greedy algorithm, defined in Section 5.1.1. The two models were evaluated on how many orders were met, followed by which model gave a lower operation cost.

5.1.1 Baseline Algorithm (Greedy)

The baseline algorithm greedily picks which edges to electrify based on the estimated weekly demand for each O-D pair. To do so, this approach does the following:

1. Sort the O-D pairs with respect to the estimated weekly demand
2. For each O-D pair, starting with the pair with the highest demand:
   (a) Find the shortest path for the current O-D pair.
   (b) Calculate the minimum number of charging stations required to fully electrify this path. This is done by summing over the number of stations required for each edge on this path using Equation 3.1.
   (c) If the number of charging stations required is less than the remaining budget, choose this path. Otherwise, move to the next O-D pair.

Once the edges have been chosen, we then run the same two-stage optimization model described in Section 3.3 to find the maximum demand that can be met, while minimizing the costs incurred to meet this demand to solve the subsequent routing problem.

Comparison to the Baseline

To establish context for the following results, let's recall that the model optimized for the total number of orders met first, followed by minimizing operation costs. Note that we need 9 non-city charging stations to fully electrify the network in the 4 city case and 8 in the 10 city case. We generate these results for a fixed fleet size of 8 for 4 cities and 12 for 10 cities.

4 Cities In terms of the location of charging stations, the model and baseline do similarly for all budget cases for the 4 city case, except when the budget is 7, which is shown in 5.1. The model prioritizes electrifying a direct path from Houston to Dallas, while the baseline prioritizes a direct path from Houston to San Antonio.

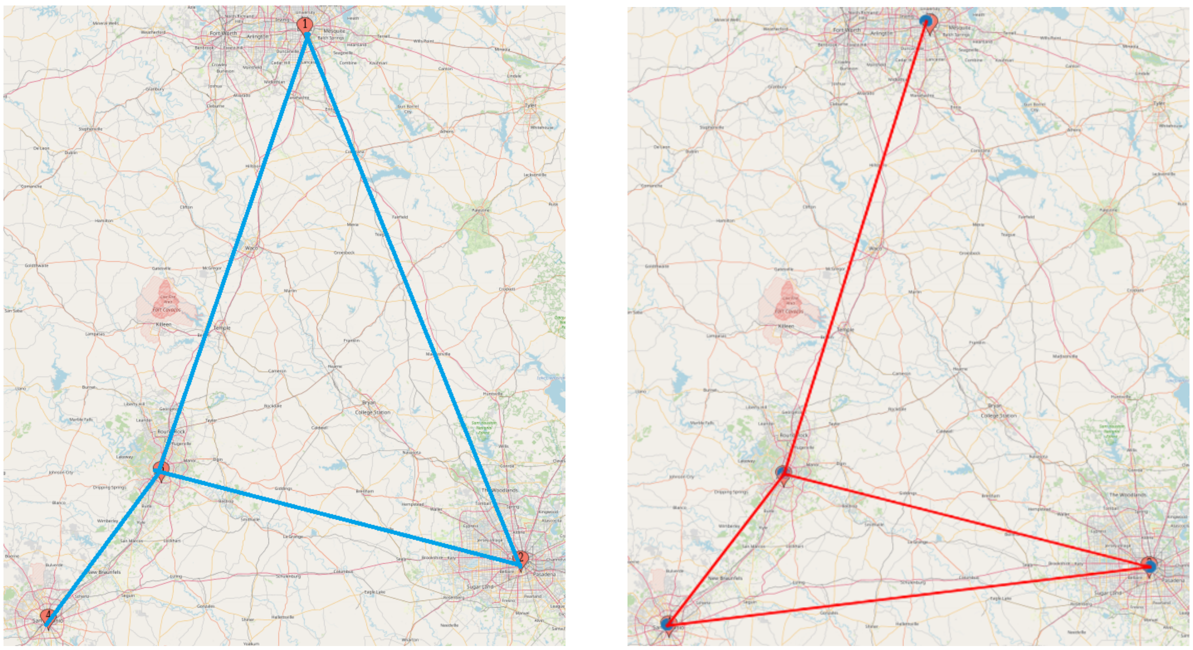

(a) Model (b) Baseline

Figure 5.1: Comparison of Model to Baseline in the 4 City Case for Budget 7

The impact of this difference is reflected in Figure 5.2b, as both the model and the baseline result in the same average costs, except for when the budget is 7. The model's choice results in a lower average cost. The alternate path from Houston to Dallas (Houston-Austin-Dallas) is significantly more expensive than the alternate path from San Antonio to Houston (San Antonio-Austin-Houston), and both of these paths have very similar traffic going along them. This longer path leads to higher costs in both travel, as well as results in more stops to charge, which means that the order takes longer to get delivered as well for EVs. The alternate is to fulfill these orders using ICEs, but ICEs cost more to drive per mile than EVs, raising costs of operations.

In terms of the percentage of orders met, both the model and the baseline perform similarly, as shown in Figure 5.2a.

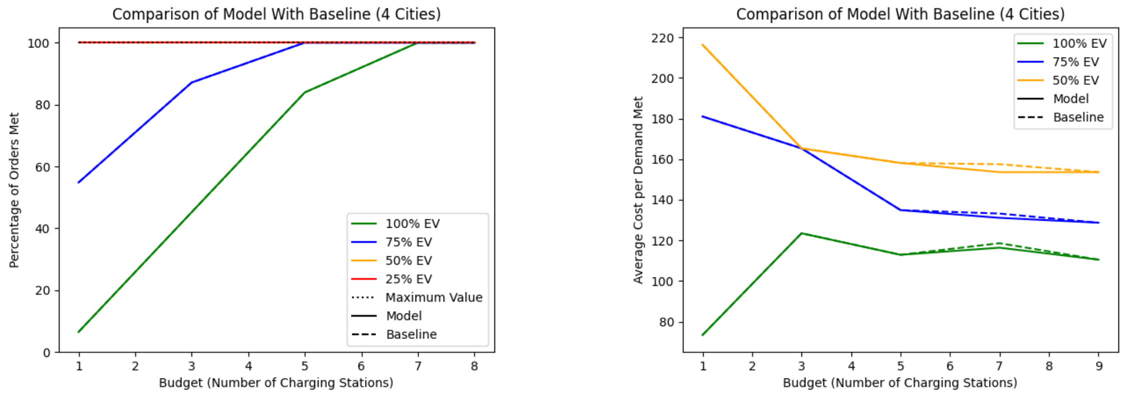

(a) Percentage Demand Met (b) Operation Costs

Figure 5.2: Comparing Model Against the Baseline for 4 Cities

10 Cities In terms of the location of charging stations, the model and baseline perform very differently for all budget cases for the 10 city case, except when the budget is 1 and 8. These differences are shown in Figures 5.3 and 5.4.

As we can see, the baseline connects paths between San Antonio and Dallas first, while the model chooses to connect San Antonio and Houston first.

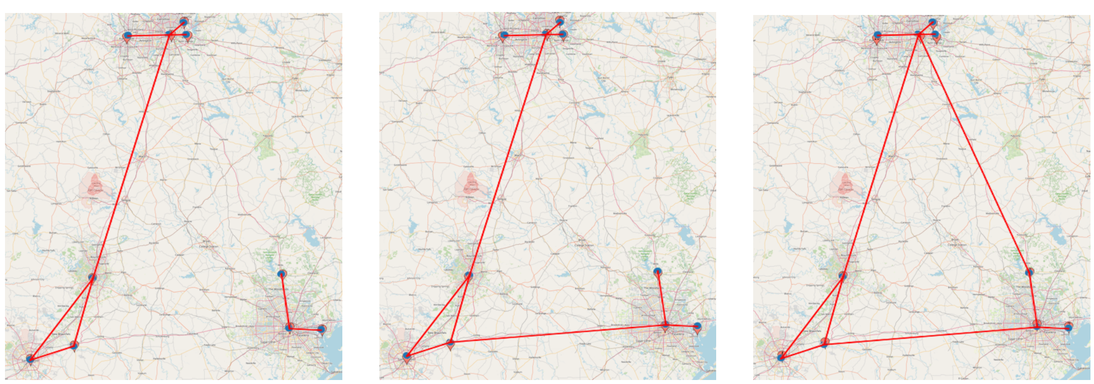

(a) Budget 3 (b) Budget 5 (c) Budget 7

Figure 5.3: Baseline Electrification of the Texas Triangle as Budget Increases in the 10 City Case

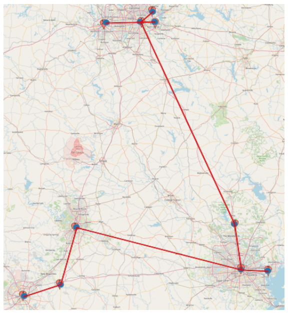

(a) Budget 3

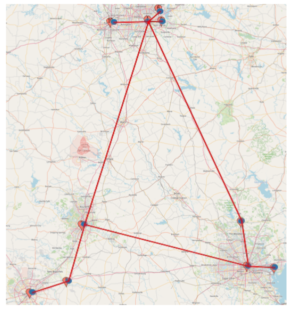

(b) Budget 5 (Majority EV)

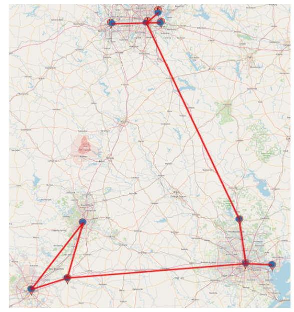

(c) Budget 5 (Majority ICE)

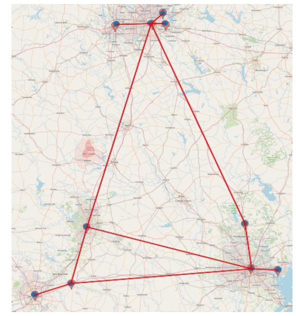

(d) Budget 7

Figure 5.4: Model Electrification of the Texas Triangle as Budget Increases in the 10 City Case

While the San Antonio to Dallas O-D pair has the highest traffic, the same is not true for the edges connecting these two cities. As a result, we can see in Figure 5.5a that there are some instances where the model can meet all the orders, but the baseline does not. When the budget is 3, for instance, all the cities are connected by the model, but not by the baseline. As a result, the model is still able to meet all the demand (albeit at a higher cost), while the baseline cannot, since it does not have enough ICEs to fulfill orders over the unelectrified paths.

Similarly, we can see in Figure 5.5b that the model generally results in a lower average cost for the reasons mentioned earlier in the 4 city case.

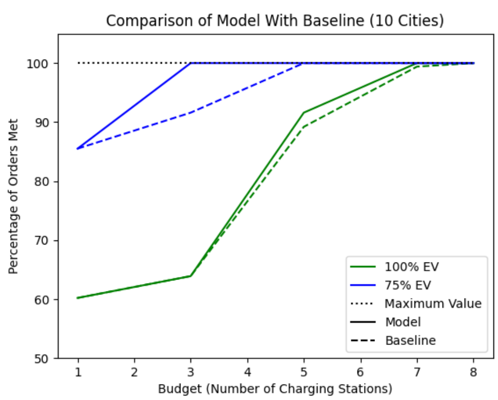

(a) Percentage Demand Met

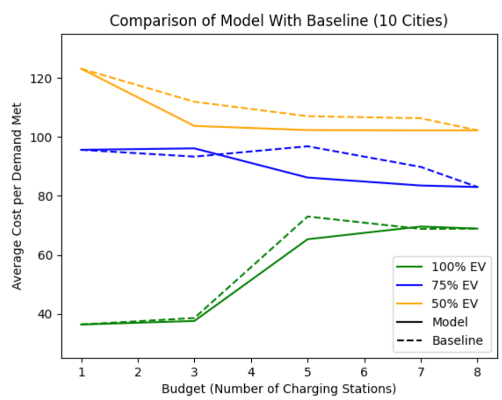

(b) Operation Costs

Figure 5.5: Comparing Model Against the Baseline for 10 Cities

Based on these results, we can see that our model tends to do at least as well as the baseline, and in fact, frequently outperforms it.

5.2 Practical Takeaways

To get better insight into the following results, we first take a look at the flow of traffic between the different cities in both, the 4 and the 10 city cases as shown in Figure 5.6. The thickness of the lines represents how much traffic flows between the two endpoints in the case where we meet all of the demand.

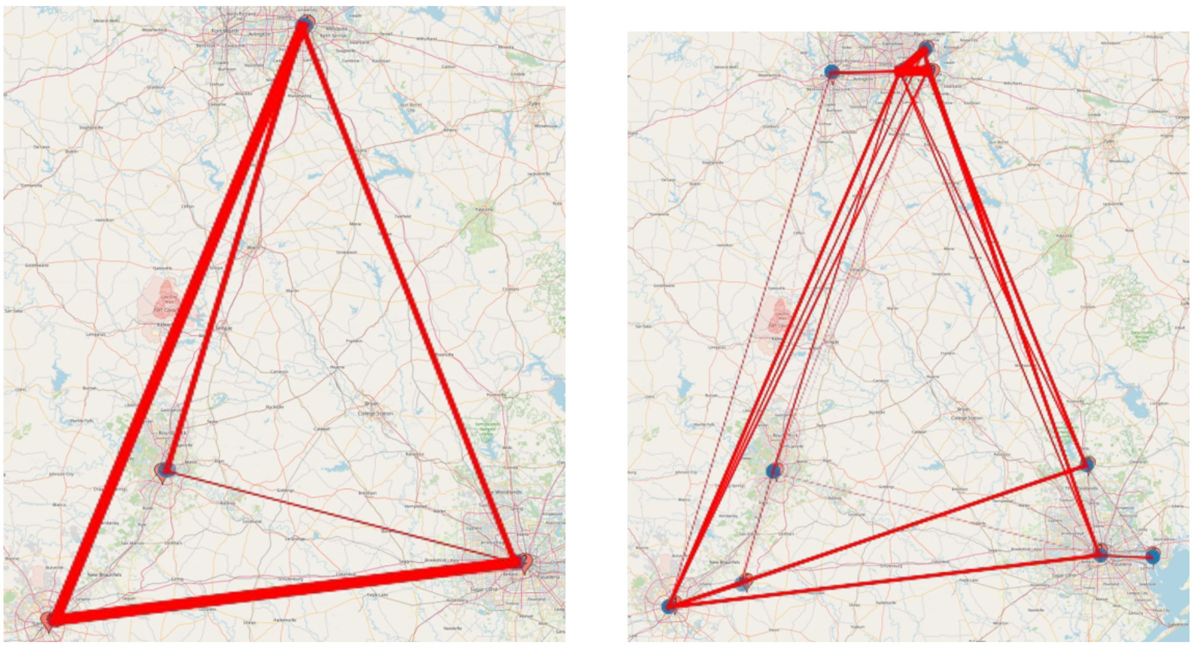

(a) 4 Cities (b) 10 Cities

Figure 5.6: Flow of Traffic Between the Different Cities

Now, we will look at how the model electrifies the network in the following three cases as we increase the budget for charging stations. Unless explicitly stated, the parameters are fixed as before. In addition, we also vary the composition of the fleet.

4 Cities For the 4 city case, figure 5.7 shows how the model goes about electrifying the network as the budget rises. We notice that the composition of the electric fleet does not impact which edges are chosen to be electrified. When we compare 5.7 to 5.6a, we notice that the model prioritizes the arcs carrying the most flow (so, the Houston to San Antonio, and San Antonio to Dallas paths are electrified as soon as the budget constraint allows it, as we can see in Figures 5.7b and 5.7c).

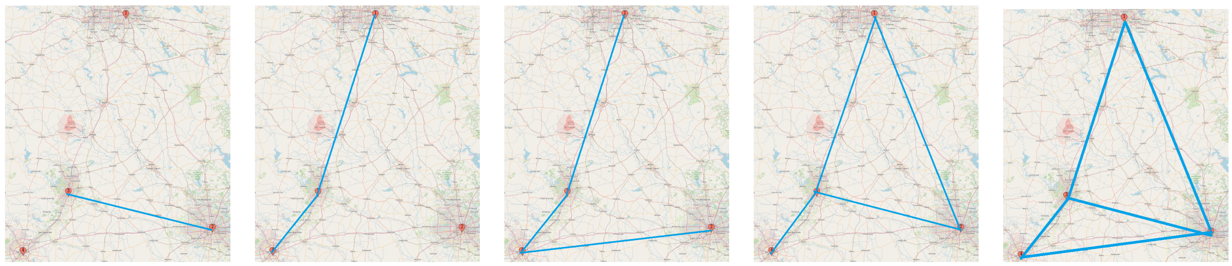

(a) Budget 1 (b) Budget 3 (c) Budget 5 (d) Budget 7 (e) Budget 9

Figure 5.7: Electrification of the Texas Triangle as Budget Increases in the 4 City Case

10 Cities The 10 city case is a bit more interesting than the 4 city case, as there are more connections to make, so we can develop more insights into how the model decides on the arcs it picks, as well as confirm some of our takeaways from the 4 city case. In general, the edges picked were pretty similar across the composition of our trucking fleet. Figure 5.8 shows how the model goes about electrifying the network as the budget rises in the fleet is majority electric ($> 50\%$ EVs).

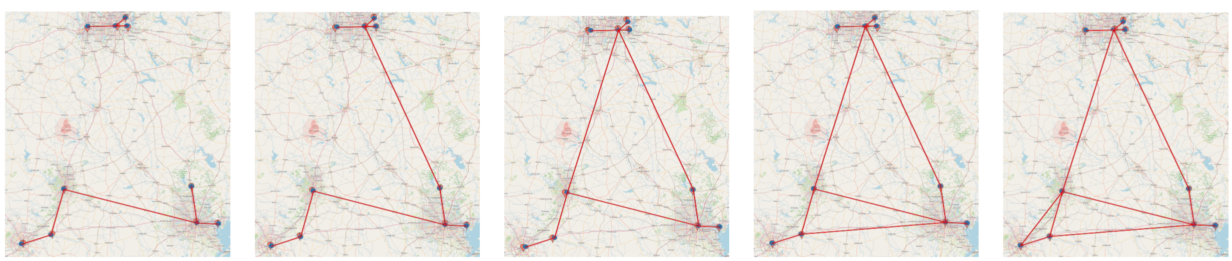

(a) Budget 1 (b) Budget 3 (c) Budget 5 (d) Budget 7 (e) Budget 8

Figure 5.8: Electrification as Budget Increases in the 10 City Case in a Majority EV Fleet

We did, however, notice some differences in the edges chosen when our budget was 5 between the figures 5.9a and 5.9b. Note that in both of these cases, all of the demand was met.

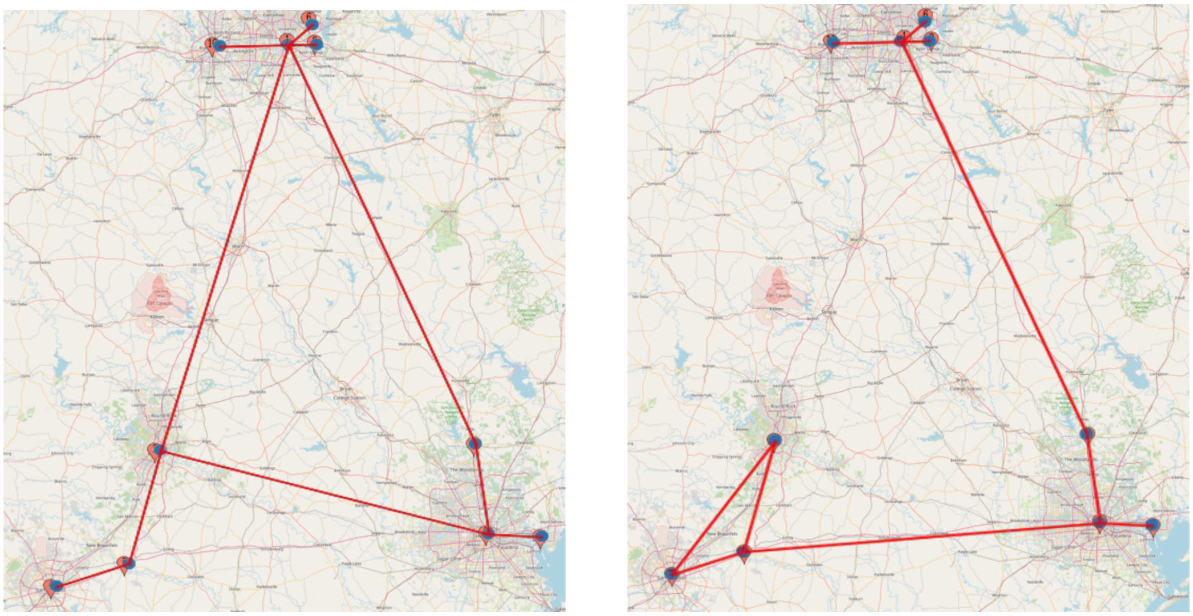

(a) Majority Electric Vehicles (b) Majority ICEs

Figure 5.9: Electrification of the 10 City Case with a Budget of 5

In Figure 5.9, we notice that in the majority EV case, the model tries to connect the entire network first, so the orders from San Antonio to Dallas (and vice versa) can be served in the shortest way possible. In the majority-ICE case, however, the model chooses to fully connect the lower half of the Texas Triangle with Electric Vehicles.

From Figure 5.6b, we know that both the San Antonio-Dallas (SA-D) route, as well as the San Antonio-Houston (SA-H) route, have a lot of traffic. In Figure 5.9, we see the model making a decision; electrify SA-D, and have the option to direct orders here over EVs, or serve SA-D using ICEs, and spend the extra budget to lower EV costs on the SA-H route.

When we have more electric vehicles (Figure 5.9a), the model chooses to electrify the SA-D route. The alternative option (Figure 5.9b) has two main flaws which makes this a better option in this case. Firstly, when we have a majority electrified fleet, we do not have enough ICEs to fully meet the demand going over the SA-D route, so this necessitates routing some of the traffic over EVs. EVs in the alternative option, however, cannot travel from San Antonio to Dallas. The path is a lot longer (San Antonio-Houston-Conroe-Dallas), and for most electric trucks, it is infeasible given a realistic daily time horizon.

In a majority ICE fleet (Figure 5.9b), we have enough ICEs to meet all the demand on the SA-D route, so it is not necessary to direct any traffic here using EVs. So, the model can choose the option that lowers the costs. Given our current demand, the model evaluated that the benefits of directly connecting SA-H using EVs outweigh the extra costs of fulfilling SA-D over ICEs. As a result, it chose to save on EV costs in the lower half of the Texas Triangle.

5.2.1 Insights

Based on the results obtained from the model, we can develop the following insights:

- The model, in general, tries to connect the network in such a way that the higher flow paths have the lowest optimal cost. Since electrical vehicles tend to have a lower operation cost than ICEs, the model tries to electrify the high-volume paths as a priority, so these can be served at the lowest cost.

- If we have enough electric vehicles to serve the entire network, the model tries to partially electrify the entire network over fully electrifying part of the network. If connecting the entire network results in a path using EVs that costs more than the shortest path using ICEs, the model will focus on fully electrifying part of the network. This is evident in how the model chooses edges when we have a budget of 5 in the 50% electric fleet case when compared to the 100% fleet case.

- Figure 5.10 highlights the importance of a gradual transition to an electrified fleet. Fully investing in EVs at a lower budget for charging stations leads to a lot of orders being unfulfilled (as bad as about 60% and 40% unmet demand in the 4 and 10 city cases, respectively). Fleets with a higher percentage of ICEs, on the other hand, can still fully meet the demand. They use EVs on electrified routes (since they have a lower variable cost), and ICEs on routes that are either not electrified yet, or are too long (such as San Antonio to Dallas in Figure 5.9b).

  As a result, figuring out how to optimally split our investments between charging stations and electric vehicles is another important decision and one that depends a lot on our needs and the trade-offs we're willing to make.

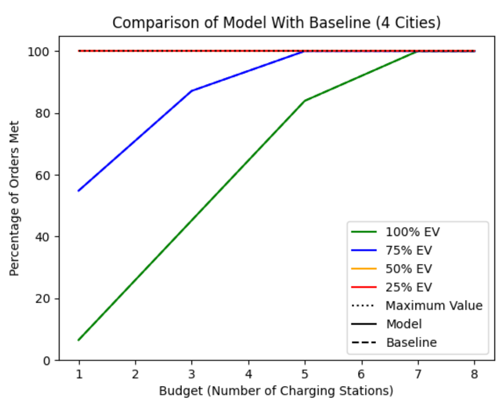

(a) 4 Cities

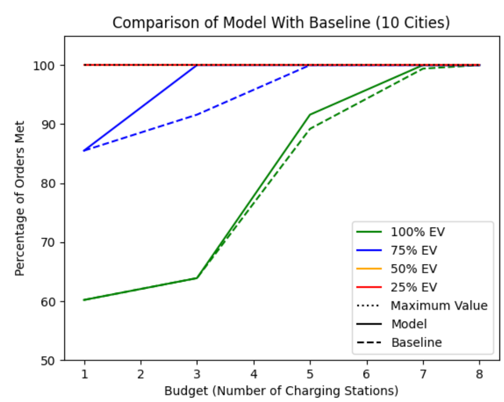

(b) 10 Cities

Figure 5.10: Percentage of Demand Met as Budget Varies

- And lastly, electric vehicles require longer operational hours to fulfill the same amount of orders as ICE vehicles. This is because EVs spend more time charging at their stations than ICEs have to wait for refueling at gas stations. So, one other takeaway from this research is that we might need to increase the size of the fleet to serve the same amount of orders, or have longer working hours, as we transition to a more electrified fleet.

These results seem intuitive and sensible, and hence, provide a good sanity check for our model.

Chapter 6

Conclusion

In this thesis, we have proposed a model that will help in the electrification of fleets. We generate our results using this model, evaluate these results, and provide practical insights into what the results mean. While the approach here is still very experimental, it does lay the groundwork for future research in this area.

6.1 Future Work/Next Steps

There are some next steps and areas for potential improvements, as well as areas for more experimentation with the model presented here.

The model currently provides the itinerary for the trucks as an approximate packing problem. There needs to be post-processing to convert the routes assigned to each truck to be more practical.

The model was designed, keeping certain common acceleration strategies (namely Bender's Decomposition) for MILPs. As a result, there is room for more experimentation using these to see how both the practical computation time, as well as the optimality gap, can be improved, and potentially provide better results.

The model also allows for experimentation with different long-term and short-term strategies simply by tweaking the objective function to a function we care about more. Electrification of the fleet, obviously takes a lot more upfront investment, while the benefits, both in terms of cost and environmental, are long-term. Finding a balance between the current capital investment and future gains is a key step that needs to be investigated. The model presented lays the foundation for any research conducted in these avenues in the future.

In addition, we can also think about potential cost savings by making use of economies of scale and encouraging collaboration among different logistics suppliers. Larger charging stations might end up being lower in cost per charging station, which might incentivize the sharing of charging stations, and hence costs, between the different suppliers.

Appendix A

Full Results

A.1 Electrification of the Texas Triangle by the Baseline

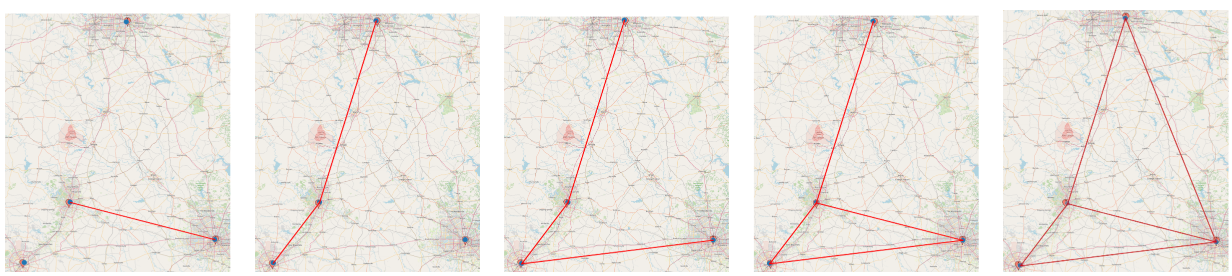

(a) Budget 1 (b) Budget 3 (c) Budget 5 (d) Budget 7 (e) Budget 9

Figure A.1: Electrification of the Texas Triangle as Budget Increases in the 4 City Case

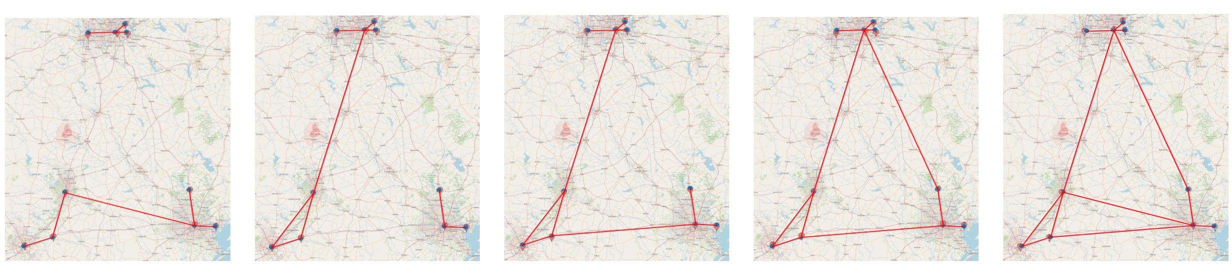

(a) Budget 1 (b) Budget 3 (c) Budget 5 (d) Budget 7 (e) Budget 8

Figure A.2: Baseline Electrification of the Texas Triangle as Budget Increases in the 10 City Case

A.2 Comparison to Baseline

For the following results, we need to remember that the model's priority is to maximize the percentage of orders met over, and then minimize costs. We notice that the model performs at least as well as the baseline in all cases. That is, the model never meets a lower percentage of orders than the baseline, and the only cases it has a higher cost is when it is meeting a higher percentage of orders, too.

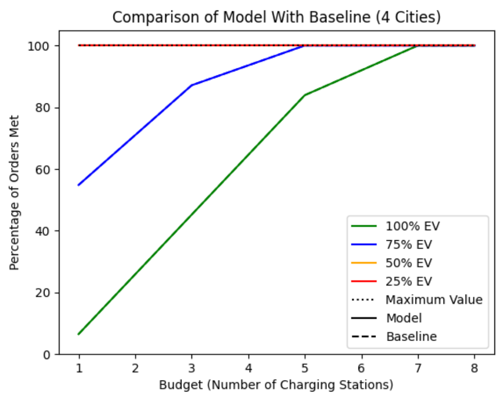

(a) Percentage Demand Met

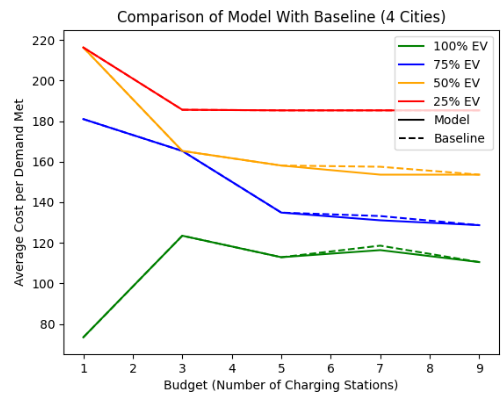

(b) Operation Costs

Figure A.3: All Results, 4 Cities

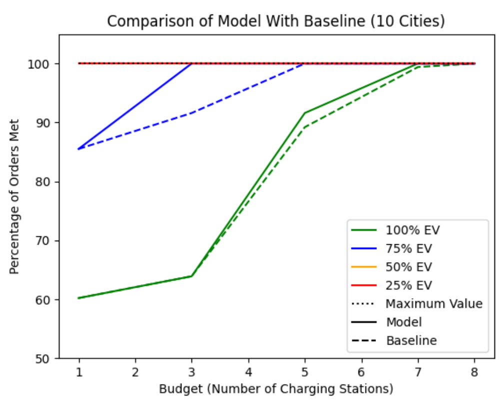

(a) Percentage Demand Met

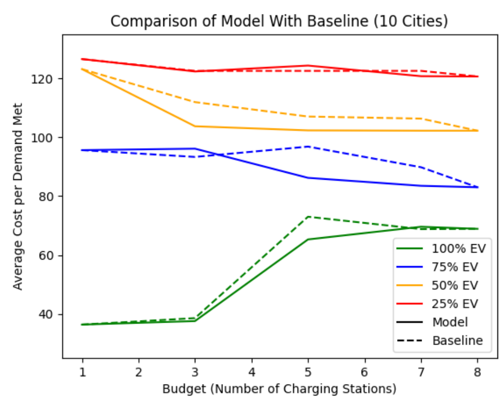

(b) Operation Costs

Figure A.4: All Results, 10 Cities

A.3 Electrification as Budget Increases

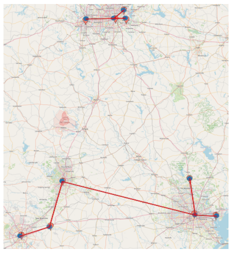

(a) Budget 1

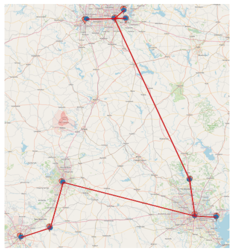

(b) Budget 3

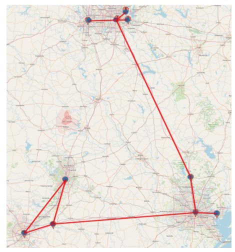

(c) Budget 5

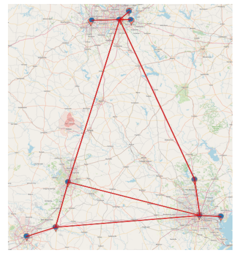

(d) Budget 7

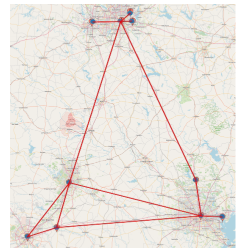

(e) Budget 8

Figure A.5: Electrification as Budget Increases in the 10 City Case in a Majority ICE Fleet

References

[1] U. E. P. Agency. "Facts facts: U.s.transportation sector greenhouse gas emissions 1990-2022." version 1.71, USEPA. (May 2024), URL: https://www.epa.gov/system/files/documents/2024-05/420f24022.pdf (visited on 06/08/2024).

[2] I. Crawford, "How much co2 is emitted by manufacturing batteries," *URl: https://climate. mit. edu/ask-mit/how-much-co2-emitted-manufacturing-b flatteries (visited on 05/19/2023)*, 2022.

[3] K. Shin. "Lifecycle emissions of electric cars vs. gasoline," Green Energy Consumers Alliance. (Jan. 2023), URL: https://blog.greenenergyconsumers.org/blog/lifecycle-emissions-of-electric-cars-vs.-gasoline (visited on 06/16/2024).

[4] T. Konstantinou and K. Gkritza, "Examining the barriers to electric truck adoption as a system: A grey-dematel approach," *Transportation Research Interdisciplinary Perspectives*, vol. 17, p. 100 746, 2023.

[5] L. Melander, C. Nyquist-Magnusson, and H. Wallström, "Drivers for and barriers to electric freight vehicle adoption in stockholm," *Transportation Research Part D: Transport and Environment*, vol. 108, p. 103 317, 2022.

[6] Business Insider. (Feb. 2023), URL: https://www.businessinsider.com/electric-trucks-longhaul-batteries-tesla-heavy-cargo-weight-problem-2023-2 (visited on 06/25/2024).

[7] U. D. of Transportation. "Comprehensive truck size and weight study, volume 2 chapter 3." (May 2022), URL: https://www.fhwa.dot.gov/reports/tswstudy/vol2-chapter3.pdf (visited on 06/25/2024).

[8] D. Nicolaides, D. Cebon, and J. Miles, "Prospects for electrification of road freight," *IEEE Systems Journal*, vol. 12, no. 2, pp. 1838–1849, 2018. DOI: 10.1109/JSYST.2017.2691408.

[9] M. Jaller, C. Otero-Palencia, and A. Pahwa, "Automation, electrification, and shared mobility in urban freight: Opportunities and challenges," *Transportation Research Procedia*, vol. 46, pp. 13–20, 2020.

[10] R. Z. Farahani, E. Miandoabchi, W. Y. Szeto, and H. Rashidi, "A review of urban transportation network design problems," *European journal of operational research*, vol. 229, no. 2, pp. 281–302, 2013.

[11] S. V. Ukkusuri and S. T. Waller, "Linear programming models for the user and system optimal dynamic network design problem: Formulations, comparisons and extensions," *Networks and Spatial Economics*, vol. 8, pp. 383–406, 2008.

[12] S. V. Ukkusuri, T. V. Mathew, and S. T. Waller, "Robust transportation network design under demand uncertainty," *Computer-Aided Civil and Infrastructure Engineering*, vol. 22, no. 1, pp. 6–18, 2007.

[13] P. Mrazovic, E. Eser, H. Ferhatosmanoglu, J. L. Larriba-Pey, and M. Matskin, "Multi-vehicle route planning for efficient urban freight transport," in *2018 International Conference on Intelligent Systems (IS)*, IEEE, 2018, pp. 744–753.

[14] F. Morlock, B. Rolle, M. Bauer, and O. Sawodny, "Time optimal routing of electric vehicles under consideration of available charging infrastructure and a detailed consumption model," *IEEE Transactions on Intelligent Transportation Systems*, vol. 21, no. 12, pp. 5123–5135, 2019.

[15] N. Chehrazi, J. Jiang, A. Wunderlich, and J. Uzoma, "The future is electric," 2022.

[16] C. Ledna, M. Muratori, A. Yip, P. Jadun, and C. Hoehne, "Decarbonizing medium-& heavy-duty on-road vehicles: Zero-emission vehicles cost analysis," National Renewable Energy Lab.(NREL), Golden, CO (United States), Tech. Rep., 2022.

[17] S. Agrawal. "Fact sheet | the future of the trucking industry: Electric semi-trucks (2023)." (May 2023), URL: https://www.eesi.org/papers/view/fact-sheet-the-future-of-the-trucking-industry-electric-semi-trucks-2023 (visited on 06/18/2024).

[18] O. R. Musin, "Properties of the delaunay triangulation," in *Proceedings of the thirteenth annual symposium on Computational geometry*, 1997, pp. 424–426.

[19] S. Suri. "Delaunay triangulation." (Oct. 2019), URL: https://sites.cs.ucsb.edu/~suri/cs235/Delaunay (visited on 06/18/2024).

[20] American Automobile Association. (Jun. 2024), URL: https://gasprices.aaa.com/?state=TX (visited on 06/21/2024).

[21] V. Khamesra, AtoB. (Nov. 2022), URL: https://www.atob.com/blog/average-mpg-for-a-semi-truck (visited on 06/21/2024).

[22] Electricity Plans. (Apr. 2024), URL: https://electricityplans.com/average-electricity-bill-in-texas/ (visited on 06/21/2024).

[23] G. Leach, "Improving worst-case optimal delaunay triangulation algorithms," in *4th Canadian Conference on Computational Geometry*, Citeseer, vol. 2, 1992, p. 15.