Modelling the Impacts of Rebalancing Strategies on Bike Share Toronto

Author
Spencer McNee, MASc
Transportation Planner
Arcadis IBI Group

Co-Author
Prof. Eric J. Miller, PhD
Director, University of Toronto Transportation Research Institute
University of Toronto

Prepared for Applications of Data in Transportation Planning of the 2023 TAC Conference & Exhibition, Ottawa, ON

Abstract

Bike Share Toronto is a docked bike share system ("System") that operates within the City of Toronto. It began operating in 2011 and has expanded to include 625 stations as of June 2022. This paper uses a microsimulation model of the System to examine the operational challenge of rebalancing bike share networks. Using simulations of the System's operation each day in 2021, this paper compares the impact of three different rebalancing scenarios upon rebalancing operations in the system.

The three scenarios cover: (a) the "as is" condition using observed rebalancing operations; (b) a worst-case scenario where no rebalancing operations are conducted, and (c) an optimized scenario where rebalancing operations are planned with perfect knowledge of ridership patterns. The optimized scenario offers a theoretical maximum efficiency to better understand how operations could be improved.

The results of the model's analysis show that the number and length of delays where a user must relocate to another station due full or empty stations decrease dramatically between the worst-case scenario and the "as is" scenario. Under the optimized scenario, users experience fewer delays and the tour lengths of the trucks performing rebalancing operations are 36% lower than in the "as is" scenario. These results highlight the potential for improved forecasting, route planning and rebalancing to reduce the System's operating costs and improve user experience.

1 Introduction

Urbanization trends around the world continue to drive strong growth in city populations. In Canada 73.7% of the population now lives in a large urban area, defined as an area with 100,000 or more people, and these urban areas are where most of the population growth has occurred from 2016 to 2021 (Statistics Canada, 2022). Improvements to cycling infrastructure and the encouragement of bike transit offer compelling solutions for cities facing the negative impacts of congestion and grappling with the challenges of decarbonization. Cycling offers users an active, inexpensive and zero emissions method of transportation. In the City of Toronto 46% of trips are taken by car, compared with 13% made by foot or bike. However, many city trips currently completed by cars are short, the median trip length by car is 5.5 km, making them strong candidates for completion by bike instead (Bess Ashby, 2018). To maximize cycling uptake, cities can encourage bike trips by providing safe cycling infrastructure, which increases people's willingness to cycle (El-Assi et al., 2017), and by offering bike share systems. These systems lower the barriers to taking up cycling by removing storage requirements and mitigating the risk of stolen bikes. Bike share systems also facilitate both one-way trips by bike and trips paired with public transit, providing a convenient last mile solution for transit journeys (Kim & Cho, 2021).

Traditional transit activity models focus on modelling distinct walking, transit, auto and cycling modes. However, the recent progression of technology has given rise to a new generation of mobility services that fall along a spectrum of use. The flexibility offered by these services provides interesting new challenges for modelling transportation systems. Modelling the interaction between the system operator and the user becomes key to model utility and accuracy. (Calderón & Miller, 2020).

Bike sharing systems are an example of a growing flexible mobility service whose users operate along this spectrum. The City of Toronto's own docked bike share system, Bike Share Toronto, began in 2011 with 80 stations and 1,000 bikes (*Need to Borrow a Bicycle?*, 2011) and has grown to a system of 625 stations and over 6,000 bikes in 2022 (Toronto Parking Authority, 2022). One of the main challenges for bike sharing systems is ensuring that the system remains in balance, meaning that there are available bikes and docking stations when users start and end their journeys. To achieve this, operators must "rebalance" the system by moving bikes between stations to manage the supply. This paper uses a microsimulation model of the System to examine the operational challenge of rebalancing bike share networks. Using simulations of the System's operation each day in 2021, this paper compares the impact of three different rebalancing scenarios upon rebalancing operations in the system.

2 Literature Review

2.1 History of Bike Share Systems

Bike share systems have exploded in popularity in recent years and many cities around the world now have extensive bike share networks. The history of bike share systems dates to 1965, when Amsterdam instituted a small system of "white bikes" meant for shared use. Ultimately, the system lacked effective security and the bikes were regularly vandalized and stolen (Fishman, 2020).

Starting in the mid-2000s systems started to arise which used docking stations for bikes and featured payment methods tied to memberships and credit cards (Fishman, 2020). These systems improved upon the security issues that beset previous bike share system attempts and the systems grew rapidly. Many such systems continue to successfully operate in a wide range of cities around the world today. The system studied in this paper, Bike Share Toronto, is such a system.

Bike Share Toronto is a docked bike sharing system in the City of Toronto. It operates using the bike and docking system developed in Montreal for the Bixi bikeshare system. The Bixi platform is used in many bike share systems including those in New York City and London (*Smart Bike-Sharing Systems for Cities*, n.d.). The system in Toronto is owned by the Toronto Parking Authority, a municipally owned corporation that subcontracts the operation of the system to the mobility operations company Shift Transit.

Bike Share Toronto was originally launched in 2011 under the name Bixi with 80 stations and 1000 bicycles (*Need to Borrow a Bicycle?*, 2011). The system was originally owned by the Public Bike System Company ("PBSC"), who created the Bixi bike share system in Montreal, but was taken over by the Toronto Parking Authority in 2014 after PBSC filed for bankruptcy. The system was renamed Bike Share Toronto after it was taken over (*New Name, Look and Prices for Toronto's Bixi*, 2014). Since the Toronto Parking Authority assumed control of the system, it has grown to over 600 stations and over 6000 bicycles.

2.2 Predicting Bike Share Demand

Many studies of bike sharing systems attempting to predict demand focus on overall ridership trends and provide estimates for daily or yearly ridership numbers. These are necessary for long term planning system expansion and infrastructure expansion but do not provide sufficient granular detail to model the daily activity in a bike share network for rebalancing planning.

When predicting demand, factors including weather and the time of day, week, and year must be taken into account. (Ashqar et al., 2019). The following models have analysed other factors that influence bike share demand.

A study of the bike share system in the San Francisco Bay Area compared several dynamic linear models to predict bike volumes at stations. These models used snapshots of the bike system as input data. They performed well in predicting bike volumes up to two hours in the future. Large spikes in station activity, such as those caused by rush hour, resulted in higher estimation errors (Almannaa et al., 2020).

Deep learning methods have also been employed to provide short term forecasting of bike demand. A study used data inputs from bike share systems in Italy and employed Bidirectional Long Short-Term Memory networks to predict trips based on historical weather and ridership observations (E. Collini et al., 2021). Another study combined two levels of neural networks to capture the spatial and temporal correlations between stations (X. Yang et al., 2020).

Studying the clustering of stations has also been an effective tool to predict bike station demand. A study in Washington, DC found that groups of stations in a geographic area had a much smoother and more predictable ridership demand. The study improved on models that focus on predicting demand at specific stations by clustering stations located in a geographic area and then applying machine learning algorithms to predict demand at the station cluster (A. A. Ramesh et al., 2021). Using clustering to predict cycling ridership is effective because some of the factors affecting ridership within a given area, such as weather, will impact all stations within that area evenly. Studies have successfully evaluated cycling routes using a clustering approach to predict link level ridership (Beitel et al., 2017).

2.3 Rebalancing Strategies

One of the main challenges of operating a bike share network arises from dealing with the problem of rebalancing available bikes at docks across the system. Without rebalancing, the system can quickly become less useful as people are unable to start or finish trips at their desired locations. Operators of bike share systems must work to ensure that full or empty stations do not prevent trips from occurring. To maintain balance within a bike share system, the operator must run rebalancing operations to move bikes between stations and enable the efficient operation of the network.

Different system operators may have distinct network operation and rebalancing objectives. For example, some operators identify and prioritize the operation of critical stations, such as those near transit stations, while others favor a more balanced approach across the network (Médard de Chardon et al., 2016). Cities have also used different strategies to incentivize rebalancing. For example, in New York City the bike share operator uses a rewards program that encourages uses to make trips which rebalance the network (*Bike Angels | Citi Bike NYC*, n.d.). The success of a network's rebalancing strategies can have important impacts on users' perceptions of the bike share system's reliability and utility. The significant impact of a successful rebalancing strategy is important for system owners and stakeholders invested in the success of the bike share system. There have been several studies which have studied this important issue using different prediction horizons and modelling techniques**.** These have included dynamic forecasting of station occupancy, as well as different decision making and vehicle routing strategies (Brinkmann, 2020).

For a rebalancing strategy to be an effective tool for decision making in practice, a model must be sufficiently computationally efficient to provide a system operator with answers quickly and consistently. Models have sought to meet this goal using different strategies and formulations. Models studying this issue can be broadly categorized into two groups. Static rebalancing models rely on a pre-determined set of rebalancing movements across a given time period, while a dynamic model can vary in strategy throughout the day.

One example of a static rebalancing model is Ren et al's evaluation of mixed integer linear programming formulations of rebalancing operations (Ren et al., 2020). This analysis did find optimized truck routings for rebalancing bikes, but it formulated the problem with all rebalancing operations conducted overnight, which did not capture the evolution of the system throughout the day. This also differs from the operational pattern in the City of Toronto, where the system operator rebalances the system during the day as well as overnight.

Focusing on the optimal rebalancing strategy is not the only consideration for the rebalancing problem. Predicting demand effectively so that operators understand when stations will be full or empty is a required first step before applying a rebalancing strategy. Dynamic rebalancing models work under the assumption that demand prediction and a subsequent rebalancing strategy can evolve as new information is available. One study of the bike share system in New York City provides an example of dynamic formulation. Researchers dynamically predicted when stations would next be full or empty based on a birth-death process. This "time until unavailable" was then used to prioritize station rebalancing and a heuristic vehicle routing algorithm applied to plan truck trips. The efficacy of this approach was measured by the fraction of that time stations were out of service, the number of rebalancing operations, and the distance travelled by trucks (Chiariotti et al., 2018).

Once predictions of ridership have been made, solving the vehicle routing of rebalancing operations on a bike share network in the most computationally optimal way is also computationally expensive. This limitation has led several studies to propose different heuristics for solving the rebalancing issue. Reducing the choice set, as well as using heuristic solutions instead of definitively optimal ones, meaningfully reduces the computation time for solving the problem, allowing solutions to be generated within a timeframe that could be useful to an operator if it were to be operationalized in the real world. One proposed approach to this problem has been to first predict the bounds of occupancy at each bike share station that will lead to meeting all demand, then to cluster stations into self-sufficient clusters that can all rebalance each other, and finally to optimize a single vehicle routing problem ("VRP") (Schuijbroek et al., 2017).

Simulation models offer a platform to compare different approaches to rebalancing a bike share system. Jin et al implemented a simulation framework testing data from the bike share system in New York City. This study provided a simulation framework of the bike share system operation, using a non-homogenous Poisson process for simulating station demand and a discrete distribution to determine destination. Travel times were modelled by a shifted Gamma distribution. The system also modelled the rate of bike deterioration leading to removal from stations for repair. Finally, the study optimized rebalancing strategies to minimize truck travel time and unloading time (Jin et al., 2022).

3 Data

This paper uses data from the City of Toronto's Open Data Portal. This data source offers three types of data about the Bike Share Toronto network. Together, these data were used to create the model discussed in this paper and they provide the basis for the following analysis. The first type of data provided by the portal is trip data. The portal makes available a log of every bicycle trip made on the Toronto Bike Share Network with details such as the start time, the end time, the start station, and the end station. The second type of data available is the station information. This includes the location, unique identifiers, names, and capacities for each station in the network. The third type of data provided is real time details about the state of each station on the network. This data is only provided in real time, it is not logged. In order to capture this third type of data and understand the progression of the Bike Share Toronto network over time, a script on Google Scripts was used to collect snapshots of the system every minute from January 2021 to July 2022.

In addition to these data provided by Bike Share Toronto, the analysis uses the City of Toronto's road network map. This data is also provided by the Open Data Portal for estimating distances and travel times between stations on the Bike Share Toronto network. The tracked travel modes include walking and driving by truck between bike share stations. Bike travel times between stations was estimated using the historical trip data.

4 Modelling Bike Share Toronto

The analysis in this paper is based on a micro-simulation model of the Bike Share Toronto network. It has been coded in Python and it utilizes the Pandas and NumPy libraries. The bike share stations are represented in a stations object that tracks the current number of bicycles docked at each station, as well as information about the station including station identifiers, locations, and its total number of docking points. The world object tracks the current time in the simulation, the travel time matrices between stations, the stations object, and the trips currently in progress. When a trip is registered in the system, the world object applies the trip by removing a bicycle from the station at the start point of the trip and logs the trip in an array which contains the expected arrivals using the stored travel time matrix stores.

When time is incremented in the simulation, the trip list is evaluated for the time increment, and the trips are applied to the world. This functions by subtracting the departures from the appropriate stations and logging the anticipated arrivals in the expected arrivals matrix. Next, the expected arrivals matrix is evaluated. Any arrivals within the time increment are added to the appropriate stations and the matrix is recentered to have the current time in the matrix's starting position.

It is important to warm up the model when simulating a day or any other time period. This involves applying trips for a period before the desired analysis so that there are trips in the system stored in the expected arrivals matrix in addition to the bicycles in docks tracked in the stations objects. To achieve this, the model also includes a method to rebase the stations in the world to a desired level of occupancy while keeping the expected arrivals matrix. A standard two hour warm up period was applied for the purposes of the analysis. This period was deemed sufficient to accurately represent ongoing trips, as 99.2% of all trips are shorter than 120 minutes.

4.1 Simulating a Day's Operation

The simulation of a day uses real trip data from observed trips on the network. These trips give a travel time that is used in the simulation. However, for the simulation of different situations and travel scenarios, when an observed trip length is not available, travel time matrices between each pair of stations on the Bike Share Toronto network for each mode are used. For walking and truck travel times, the shortest path between each pair of stations is computed using the City of Toronto road network, with the walking matrix ignoring road directionality and the truck matrix obeying road directions. An average travel speed of 5 km/h is then applied to the walking distances and a speed of 30 km/h is applied to the truck travel distances. Historical trip data is used for the bike travel time matrix. The median travel time between a pair of stations is taken for use in the simulation. However, some pairs of stations do not have a sufficient number of trips between the origin destination pair to provide an accurate median. In such cases, the shortest path distance and an average travel speed of 12 km/h are used to generate the travel time. Although there are examples of trips starting and ending at the same station with a wide range of travel times, these trips are considered to be recreational trips and are difficult to recreate in a simulation setting. For the purposes of the micro-simulation model, if there is no observed trip to draw from, a trip starting and ending at the same station is considered to have a travel time of zero.

Beyond the measures discussed above, the behaviour of users encountering full or empty stations needs to be simulated in order to simulate the functioning of the bike share network. The station object compares its updated occupancy with its capacity when trips are applied. If the new occupancy is greater than the station's capacity, trips above the capacity must be rerouted. When a trip is rerouted because a station is full, the model uses the bike travel time matrix and finds the closest station by travel time that is not at capacity. The expected arrivals matrix is then updated with a new arrival at this nearby station at a time equal to the travel time in the future. If the new occupancy is less than zero, then the trip must also be rerouted. For this empty station case, the model uses the walking travel time matrix and finds the closest station by travel time that is not empty. The expected arrivals matrix is then updated with a departure at a time equal to the walking time in the future and an arrival at the original destination at a bike travel time between the new origin and original destination after.

This behaviour for full and empty stations is considered valid because of the smartphone app which gives users real-time occupancy information on the network. A user who encounters a full or empty stations could reasonably check this app to confirm where the nearest station with available bikes or docks and travel there instead of simply travelling to the nearest station.

When rerouting trips, the system logs both the number of trips delayed and the length of delay. For trips rerouted by a full station, the delay is counted as the travel time from the full station to the new destination. For trips rerouted by an empty station, the delay is counted as the walking time to the new origin, plus the travel time to the destination, minus the original travel time. This delay is always considered a delay and if the resulting new travel time is less than the original travel time, the delay is considered to be zero, not a negative value. This is to capture the reality that a user's original plan could not be executed and the change should not be considered as an advantage even if the travel time would be shorter.

Finally, there are the movements that the system operator makes. In Bike Share Toronto's case, the operator runs trucks around the city to remove bikes from stations at or nearing full capacity and to add bikes to stations which are or approaching empty. The model represents this with a bike operator object. This object stores a list of planned truck movements and at each time increment, the model world checks the bike operator to see if there are any truck trips to initiate. A truck trip has an origin, a destination, a start time, a duration, the number of bikes removed from the origin, and the number of bikes delivered to the destination. At each timestep, the model checks if there is a new truck trip starting. If there is, the appropriate number of bikes will be removed from the station in question, and the appropriate number of bikes will be added to the expected arrivals matrix at the duration in the future. The number of trucks available and the truck routing are handled in a separate truck planning module discussed in Section 5.

Another role of the bike share system operator is to remove bikes in need of servicing from docks. When the model system is initiated, each station has a record of how many broken bikes are at each station. The capacity of the station is calculated by taking the number of docking points less the number of broken bikes occupying the station. During the simulation, it is assumed that whenever the operator performs a rebalancing operation on a station, the broken bikes are removed from the station for servicing. Once this happens the model updates the capacity of the station to reflect the newly available docking points.

The Bike Share Toronto operator can also implement "valet-stations". These stations are staffed by Bike Share Toronto operators and have extra bikes on site at a station. The operator can add or remove bikes from the station throughout the day. This effectively lifts the station's capacity constraints. Stations that experience a sufficiently large reversing trend, such as a large number of commuters arriving in the morning and departing in the evening, can reach both full and empty conditions within the same day. Such stations are considered to be good candidates for valet-stations.

There are a couple of elements that do not exist in this iteration of the model but could be implemented in the future. The first is the process of bike condition degradation. In reality, when a user encounters an issue with a bike, such as a flat tyre, they can dock the bike and flag it as needing repair. This notifies the system operator to remove the bike from service and repair it. There is no effective way to model the wear of the bikes to predict this breakdown. This makes the process by which bikes break down random. A system where bikes break at random could be implemented into the model. This was not modelled in this case because the effect was not considered to have a meaningful impact on the analysis provided or to contribute to a more useful model. Rerouted trips were considered to use the median travel time, or the shortest path distance with an assumed average speed. A more detailed travel time that draws from the distribution of the observed travel times between two stations could be implemented but was not considered to be a major issue. The approximation was considered acceptable in this case given the number of factors that can affect travel time such as fitness, route choice, and time spent adjusting the bike which are not modelled.

4.2 Modelling Assumptions

Trip length varies along each available route on the network. An assumed average travel speed of 12 km/h along the shortest path route between two stations is considered a good approximation where detailed historical trip data between two stations is not available. In modelling the system, the median trip length between two stations is used to calculate travel time where sufficient historical data are available.

The approximation that users will take the shortest path and travel at 12km/h simplifies the model system where sufficient historical data in not available. This approximation is informed by both past studies and an analysis of trips on the Bike Share Toronto network. Using GPS traces of observed trips, a previous study of cycling route choice in the City of Toronto found that, in addition to length of path, several factors affect route choice among cyclists. The study found that cyclists will accept longer travel distances in order to ride on bike facilities such as bike lanes and cycle tracks, will avoid riding on busy arterial roads, will avoid hills, and will choose routes with fewer turns. The study found these factors to be significant and found that only 22% of the observed trips matched the shortest path. It also found that there was a 10% deviation on average from the shortest cycling path (Grond, 2016). Another study conducted using GPS traces of bike share trips in Hamilton, Ontario came to similar conclusions on this topic. The study found that the route most frequently chosen between a pair of bike share stations corresponded to the shortest path in as few as 7% of station pairs, but that the average difference between the route taken and the shortest path distance was about 10% (Lu et al., 2018).

To confirm if this held for Bike Share Toronto, the shortest path distance between each pair of stations was first computed. Based on this shortest path route, an average speed was computed by dividing the route length by the duration of the trip. While there are some outliers, in general, the shortest path distance appears to be a reasonable approximation. When calculating the speed of each trip under the assumption that the shortest path distance was taken by the rider, the distribution of speed has a mean value of 11.3 km/h. This value is slightly slower than previous studies has found for average cycling speeds (Clarry et al., 2019) of around 18 km/h which indicates that users are likely taking a path slightly longer than the shortest path in addition to riding at slower speeds due to the heaviness of the bikes and the time it takes a user to adjust their bike before departing. Taken together, these results do support the use of the shortest path at an average speed of 12 km/h as a reasonable approximation of travel time where historical data in unavailable.

An analysis of trips to and from station pairs found no clear trend between the duration of a trip and the elevation gain or loss meaning the trips were of similar duration in the direction they gained elevation as that where they lost elevation. For this reason, elevation was not added as a factor in the model. Similarly, precipitation was found to lower ridership globally on the system, this making the risk of station imbalances on a short time scale less likely. Precipitation impacts are not modelled but could be by applying a reduction factor when heavy precipitation is expected.

5 Rebalancing

5.1 Detecting Real World Rebalancing Movements

One of the goals of this analysis is to compare the real world operation of the Bike Share Toronto network with different theoretical operating methods. The procedure described below outlines how the observed, real world rebalancing operations were detected.

The dataset includes snapshots of the bike share system and a list of all trips taken. The detection was applied for each day within the analysis period. The model first identifies if there is an appropriate snapshot to use as the starting point for the analysis day. The model then searches for a snapshot at least two hours before the simulation world. The warmup world is initialized and the system simulates all the trips in the trip list until the target snapshot. The occupancy of each station in the world is then set to match the target snapshot. This yields a model world that is warmed up with trips in the expected arrivals matrix and with the observed occupancies. A simulation is then conducted until the end of the day that ignores capacity constraints. Ignoring capacity constraints allows stations to have occupancies above their capacities or to have occupancies below zero. At an hourly interval, the occupancies of the simulated stations are compared with snapshots of the system. The difference between these two sets of occupancies shows where the system operator has added or removed bikes. Since some minor discrepancies can occur between the snapshots and the simulated world arising from the rounding of travel times to the nearest minute, for example, a filter is then applied to look for changes in occupancy of three bikes or greater. These jumps are then logged in a list detailing the time and amount of observed rebalancing operations.

In order to understand how well these rebalancing events work, the system simulates the day while enforcing capacity constraints and uses this generated list of observed rebalancing events as the basis for truck movements. In this way, the simulation provides details about the trips that still are delayed due to full or empty stations. In a perfect simulation this would be zero since the applied trip list shows trips that did happen on the network, making them all possible in reality. However, there still are discrepancies that result in some trips being rerouted. This effect is minor. 97.6% of all trips were completed in the capacity constrained simulation when the rebalancing truck trips were applied during the 268 days with available sufficient data to perform this analysis.

5.2 Computing an Optimal Rebalancing Strategy

The above section outlined the method used for detecting rebalancing within the network. With the benefit of perfect knowledge, meaning knowing the daily trip list in advance, can a more efficient rebalancing strategy be developed?

The goal of the optimal rebalancing strategy is to find an optimal set of rebalancing operations which would satisfy all of the observed demand for a given day. In order to achieve this, the model system is initialized and warmed up in the same manner as it is in preparation for a simulation. The start time of the simulation is considered timestep zero. For each station on the network, an array is created to log the change in occupancy in each timestep of the simulation. In this analysis, this is 24 hours or 1440 timesteps, with each timestep representing one minute. First, the outstanding trips in the world from the model warmup are logged into this trip array. Next, all trips departing from the station are subtracted from their respective timesteps and all trips arriving to the station are added to their respective timesteps. The cumulative sum of the trips array is taken, yielding the cumulative change in occupancy at the station. The occupancy of the station at the start of the simulation, timestep zero, is then added to the cumulative sum to give an array with the projected occupancy at each timestep.

This occupancy array is then checked for values below zero, or above the station's capacity. If none are found, then the station is in balance for the day and no rebalancing operations need be performed. If a violation of the capacity check is found, then rebalancing planning is triggered.

The first timestep with any capacity limit violations is evaluated. First, any broken bicycles at the station are removed, and then station capacity is updated. Then each rebalancing amount is evaluated, from removing the full capacity to adding the full capacity. If the problem timestep has an occupancy above the station capacity, the maximum rebalancing is adjusted to zero. If the occupancy is below zero, the minimum rebalancing is set to zero. Each rebalancing amount in the range is evaluated by creating a new occupancy array for the timestep onwards. If a rebalancing amount is found that will satisfy the capacity constraints until the end of the simulation, that value is selected. If no value will satisfy the constraint until the end of the simulation, the value that maintains balance for the longest is chosen. This process is then repeated until the capacity constraint can be met for the full simulation period or until a maximum of ten iterations have been performed.

The output of the algorithm is a table with the load needed to maintain balance at each station and the time window which these rebalancing operations must be performed. This load is fed into the dispatcher in order to execute these rebalancing operations from the optimal rebalancing strategy in a simulation.

5.3 Building the Truck Routing Plan

First a “naïve” truck rerouting plan is applied. This plan is not constrained by a number of trucks or optimized for efficient truck routing. This problem can then be formulated as a VRP. Solving a VRP in a closed form solution is complex and can be extremely computationally expensive. On a busy day in the Bike Share Toronto network, with 625 stations, there can be as many as 246 stations that need rebalancing. As several assumptions and conditions already exist within the formulation of the model, a heuristic solution is appropriate to apply. The Clark-Wright Savings Algorithm is a heuristic algorithm capable of providing a good solution to the VRP in a computationally efficient manor (Richard Larson & Amedeo Odoni, 1981). The object of the algorithm is to minimize the travel time of trucks performing the rebalancing operation by combining stops at different stations into tours. For computational efficiency, reducing the choice set of stations in the vehicle routing problem is key.

Step 1 of the savings algorithm computes the travel time savings that could be realized by combining each pair of stations needing rebalancing. Step 2 sorts the savings from largest to smallest. Step 3 looks at the savings in order and combines stations into a tour if no constraints are violated. The constraints include the tour position constraint, the timing constraint, and the capacity constraint. The tour position constraint stipulates one of three conditions be met: that neither station is in a tour, that only one station is in a tour and the other station is at the start or end of the tour, or that both are in tours and both are at the start or end of the tour. The timing constraint stipulates that the proposed new tour must be possible given the delivery requirements for the rebalancing operations. The truck must be able to arrive at each station before the time the rebalancing is required with each tour starting and ending at the depot. In the case where there are multiple deliveries to a single station, each subsequent delivery must be after the previous one. The capacity constraint stipulates that the truck must not exceed its capacity or have below a capacity of zero at any point along the tour, considering the collection and delivery of bikes at stations along the tour. Step 3 is repeated until all pairs have been evaluated.

The Bike Share Toronto network had 6,850 bikes in 2021 (Toronto Parking Authority, 2022). In a simulated day the total number of bikes observed on the network is less than 6,000 on most days and did not exceed 6,600 bikes in any day in 2021. For this reason, it is considered a reasonable assumption that a truck performing rebalancing operations can depart from the depot with any number of bikes and would not need to collect bikes from a station unless it was in need of rebalancing.

There is a Bike Share Toronto depot located at 25 Booth Ave in Toronto. It is assumed for the formulation of the vehicle routing problem that this is the one depot where all truck tours will start and end. Truck travel times between each pair of stations and between each station and the depot are estimated using the shortest path distance on Toronto's road network and an assumed average travel speed of 30 km per hour. QGIS was used to compute the shortest path on the network.

6 Model Experiments

The analysis examines three simulation scenarios for the Toronto Bike Share network. The first is a scenario where no rebalancing is conducted. This allows for a baseline measure with which to compare. This measure shows the number and duration of delays that would be experienced on the network without the system operator's intervention. The second scenario uses the rebalancing actions observed on the network. This scenario shows how the system is operated in reality and provides a benchmark against which other scenarios can be assessed. Finally, the third scenario seeks to posit an optimized rebalancing strategy that is theoretically the most efficient at network rebalancing.

All scenarios in this analysis use the observed trip list for the simulation. It is possible that the true demand is greater than this observed trip list. A user could arrive at a station and, finding it empty, choose to use a different mode of travel instead. In the simulations, 85% of empty station reroutings resulted in a delay of less than eight minutes. Rerouting trips may show up as other observed trips on the network where the user has chosen to walk to the nearest station. As there is no effective method or data available to quantify the number of such lost trips, the use of the observed trip list is considered to be valid for all of the simulation scenarios.

This analysis examines each day independently. Each analysis day is reset to the observed state of the system at that point. Some stations tend towards imbalance over a longer time period. For example, station 7239 experiences an average net gain of 1 bike per day and has a capacity of 20. It follows that particular station would on average need to be rebalanced once every 20 days, but this will not necessarily be captured on any individual day of simulation. These stations do required rebalancing, but they are not considered in this analysis. Note that system operators can more easily plan for these trips, as they occur over a much longer time horizon.

7 Results

The simulation of the Bike Share Toronto network was performed for each day in 2021. It was not possible to simulate certain days due to insufficient data. For the No Rebalancing scenario and then optimized rebalancing scenario, 345 days were simulated. For the observed rebalancing scenario, 268 days were simulated. The number of days is lower because the observed rebalancing algorithm relies on a greater number of observed snapshots in order to detect rebalancing events.

7.1 Overall Scenario Comparison

Figure 1 shows a plot of the number of delayed trips, either due to a full or an empty station, plotted against the total number of trips on that day. As expected, days with higher ridership numbers have a greater number of delayed trips. There is a spread in the number of delayed trips at the higher end, suggesting that different ridership patterns impact the number of delayed trips. Figures 1 and 3 show the number of delayed trips and the total delay by day of the year. Figures 2 and 4 show the number of delayed trips and the total delay versus the number of trips observed on the date in question. The total time delay and total number of trips are both significantly higher on days with higher trip counts. Higher trip counts are observed in warmer months.

The optimized rebalancing algorithm does address nearly all needed rebalancing operations while the observed rebalancing scenario still has some amount of delay. Since all trips in the simulation are from observed trips, and were thus possible in reality, the delays in the observed rebalancing scenario are a result of imprecisions in the algorithm used to detect rebalancing.

For the trips that are delayed by full stations, most are able to be rerouted effectively. Figure 5 show a histogram showing the distribution of the delays. The empty station delays have a large number which result in no delays. This is due to the fact that the station to which users rerouted had a shorter travel time, resulting in an arrival time no later than the originally predicted time. Where delays do result from empty stations, the delays do tend to be longer than those for full stations, which is to be expected since the delay includes extra walking instead of extra cycling.

Figure 5 shows the distribution of delay length observed in each scenario. All three scenarios show a similar distribution in the length of delays arising from both full stations and empty stations. Delays from full stations are generally shorter with a mean delay in all three scenarios of 3.5 minutes. Delays from empty stations are longer with a mean delay range of 4.2 to 4.5 minutes. For empty-station delays, there are a significant number of observations of zero delay. As mentioned above, this is because there are many scenarios where the additional walking to a nearby station is offset by faster travel time along the updated cycling route.

The distribution of these delays supports the assumption that the majority of delayed trips on the system are able to be rerouted. This indicates that the use of observed trips is appropriate, as trips are likely still present even if a user had to reroute.

7.2 Scenario 1 No Rebalancing

The spatial distribution of stations that most frequently required rebalancing when simulated in 2021 show that as stations approach the core of the city, their chance of being full increases. This is expected as this is also where ridership is highest and where the greatest number of destinations are located. The spatial distribution of empty stations shows a very clear cluster in the northern part of downtown as well as near certain transit stations on the Yonge subway line and near Broadview subway station.

7.3 Rebalancing Operations

This section compares the rebalancing operations between the observed rebalancing and the optimized rebalancing scenarios. Figure 6 below shows a plot of the number of station rebalances in each simulated day for both the optimized and the observed rebalancing scenarios. Figures 7 and 8 show the number of bikes removed and added in each day respectively. The optimized scenario performs better, reducing the total number of rebalancing operations conducted. On average, the optimization scenario is able to reduce the number of rebalancing movements by 36% when compared with the observed rebalancing scenario. Figure 9 below shows a breakdown of the percent reduction in rebalancing movements. While some of the observed rebalancing operations may be addressing longer term imbalances that are not captured within a single day, this reduction shows that there are significant operational efficiencies to be gained by improved trip forecasting and optimizing rebalancing strategies.

7.4 Tour Lengths

Section 7.3 described the reduction in the number of rebalancing operations in the optimized scenario when compared with the observed rebalancing. A similar reduction is noted in the truck tour length. This is significant because on the median day the total tour length for the observed rebalancing reaches 1,058 km. Optimizing the rebalancing operation can reduce this tour length, thereby decreasing the emissions from truck traffic and lowering operational costs from person hours for operators. Figure 10 below shows a breakdown of the percent reduction in tour length when applying the optimized rebalancing strategy. On average, a 36% reduction in travel distance is observed. There are some days where the tour length in the optimized case is greater than that of the observed rebalancing scenario, indicating that further improvements could be made to the optimization algorithm and truck routing.

8 Conclusions and Future Work

Bike sharing systems are a flexible mobility option which have great potential in cities. In order to reach their potential, these systems must be effectively rebalanced to ensure that the users can start and end their rides at the most convenient location. This paper used a microsimulation model to simulate the operation of Bike Share Toronto in order to evaluate the impact of rebalancing and rebalancing strategy on delay experienced by users of the system.

The analysis covers three rebalancing scenarios: (a) the “as is” condition using observed rebalancing operations; (b) a worst-case scenario where no rebalancing operations are conducted, and (c) an optimized scenario where rebalancing operations are planned with perfect knowledge of ridership patterns, offering a theoretical maximum efficiency to better understand how operations could be improved.

The results of the model’s analysis show that the number and length of delays where a user must relocate to another station due full or empty stations decrease dramatically between the worst-case scenario and the “as is” scenario. Under the optimized scenario, users experience fewer delays and the tour lengths of the trucks performing rebalancing operations are 36% lower than in the “as is” scenario. These results highlight the potential for improved forecasting, route planning and rebalancing to reduce the Bike Share Torontos’s operating costs and improve user experience.

References

A. A. Ramesh, S. P. Nagisetti, N. Sridhar, K. Avery, & D. Bein. (2021). Station-level Demand Prediction for Bike-Sharing System. *2021 IEEE 11th Annual Computing and Communication Workshop and Conference (CCWC)*, 0916–0921. https://doi.org/10.1109/CCWC51732.2021.9375958

Almannaa, M. H., Elhenawy, M., & Rakha, H. A. (2020). Dynamic linear models to predict bike availability in a bike sharing system. *International Journal of Sustainable Transportation*, *14*(3), 232–242. https://doi.org/10.1080/15568318.2019.1611976

Ashqar, H. I., Elhenawy, M., & Rakha, H. A. (2019). Modeling bike counts in a bike-sharing system considering the effect of weather conditions. *Case Studies on Transport Policy*, *7*(2), 261–268. https://doi.org/10.1016/j.cstp.2019.02.011

Beitel, D., McNee, S., & Miranda-Moreno, L. F. (2017). Quality Measure of Short-Duration Bicycle Counts. *Transportation Research Record*, *2644*(1), 64–71. https://doi.org/10.3141/2644-08

Bess Ashby. (2018). *Transportation Tomorrow Survey 2016*. Malatest.

Bike Angels | Citi Bike NYC. (n.d.). Retrieved August 22, 2022, from https://citibikenyc.com/bike-angels

Brinkmann, Jan. (2020). *Active Balancing of Bike Sharing Systems* (1st ed. 2020.). Springer International Publishing. https://doi.org/10.1007/978-3-030-35012-3

Calderón, F., & Miller, E. J. (2020). A literature review of mobility services: Definitions, modelling state-of-the-art, and key considerations for a conceptual modelling framework. *Transport Reviews*, *40*(3), 312–332. https://doi.org/10.1080/01441647.2019.1704916

Chiariotti, F., Pielli, C., Zanella, A., & Zorzi, M. (2018). A Dynamic Approach to Rebalancing Bike-Sharing Systems. *Sensors*, *18*(2). https://doi.org/10.3390/s18020512

Clarry, A., Faghih Imani, A., & Miller, E. J. (2019). Where we ride faster? Examining cycling speed using smartphone GPS data. *Sustainable Cities and Society*, *49*, 101594. https://doi.org/10.1016/j.scs.2019.101594

E. Collini, P. Nesi, & G. Pantaleo. (2021). Deep Learning for Short-Term Prediction of Available Bikes on Bike-Sharing Stations. *IEEE Access*, *9*, 124337–124347. https://doi.org/10.1109/ACCESS.2021.3110794

El-Assi, W., Salah Mahmoud, M., & Nurul Habib, K. (2017). Effects of built environment and weather on bike sharing demand: A station level analysis of commercial bike sharing in Toronto. *Transportation (Dordrecht)*, *44*(3), 589–613. https://doi.org/10.1007/s11116-015-9669-z

Fishman, E. (2020). *Bike share*. Routledge.

Grond, Kathryn. (2016). *Route Choice Modeling of Cyclists in Toronto*. Thesis (M.A.S.)--University of Toronto (Canada), 2016.

Jin, Y., Ruiz, C., & Liao, H. (2022). A simulation framework for optimizing bike rebalancing and maintenance in large-scale bike-sharing systems. *Simulation Modelling Practice and Theory*, *115*, 102422. https://doi.org/10.1016/j.simpat.2021.102422

Kim, M., & Cho, G.-H. (2021). Analysis on bike-share ridership for origin-destination pairs: Effects of public transit route characteristics and land-use patterns. *Journal of Transport Geography*, *93*, 103047. https://doi.org/10.1016/j.jtrangeo.2021.103047

Lu, W., Scott, D. M., & Dalumpines, R. (2018). Understanding bike share cyclist route choice using GPS data: Comparing dominant routes and shortest paths. *Journal of Transport Geography*, *71*, 172–181. https://doi.org/10.1016/j.jtrangeo.2018.07.012

Médard de Chardon, C., Caruso, G., & Thomas, I. (2016). Bike-share rebalancing strategies, patterns, and purpose. *Journal of Transport Geography*, *55*, 22–39. https://doi.org/10.1016/j.jtrangeo.2016.07.003

Need to borrow a bicycle? Bixi launches in May. (2011, April 26). Thestar.Com. https://www.thestar.com/news/gta/2011/04/26/need_to_borrow_a_bicycle_bixi_launches_in_may.html

New name, look and prices for Toronto's Bixi. (2014, March 30). Thestar.Com. https://www.thestar.com/news/gta/2014/03/30/new_name_look_and_prices_for_torontos_bixi.html

Ren, Y., F. Zhao, H. Jin, Z. Jiao, L. Meng, C. Zhang, & J. W. Sutherland. (2020). Rebalancing Bike Sharing Systems for Minimizing Depot Inventory and Traveling Costs. *IEEE Transactions on Intelligent Transportation Systems*, *21*(9), 3871–3882. https://doi.org/10.1109/TITS.2019.2935509

Richard Larson & Amedeo Odoni. (1981). *Urban Operations Research*. Massachusetts Institute of Technology. http://web.mit.edu/urban_or_book/www/book/index.html

Schuijbroek, J., Hampshire, R. C., & van Hoeve, W.-J. (2017). Inventory rebalancing and vehicle routing in bike sharing systems. *European Journal of Operational Research*, *257*(3), 992–1004. https://doi.org/10.1016/j.ejor.2016.08.029

Smart bike-sharing systems for cities. (n.d.). PBSC Urban Solutions. Retrieved July 26, 2022, from https://www.pbsc.com/

Statistics Canada. (2022). *Canada's large urban centres continue to grow and spread*.

Toronto Parking Authority. (2022). *BIKE SHARE TORONTO FIRST QUARTER (Q1) 2022 UPDATE*.

X. Yang, S. He, & H. Huang. (2020). Station Correlation Attention Learning for Data-driven Bike Sharing System Usage Prediction. *2020 IEEE 17th International Conference on Mobile Ad Hoc and Sensor Systems (MASS)*, 640–648. https://doi.org/10.1109/MASS50613.2020.00083

Figures

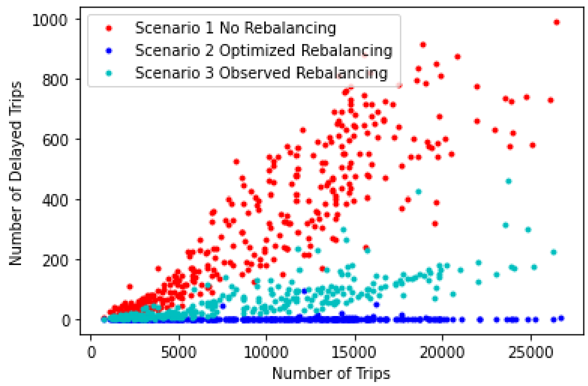

Figure 1: Plot of Total Number of Trips vs Number of Delayed Trips in a Day

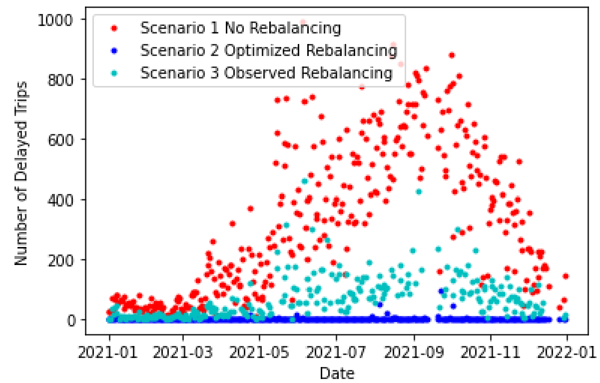

Figure 2: Plot of Number of Delayed Trips by Date

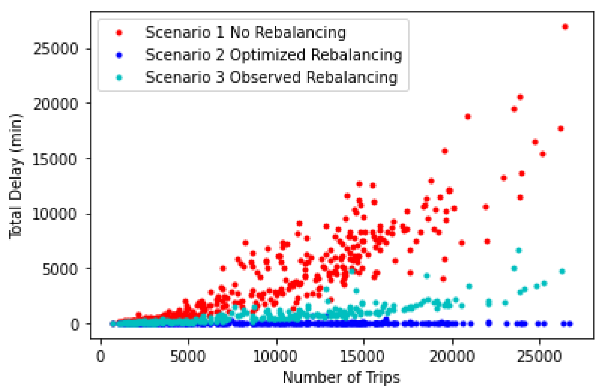

Figure 3: Plot of Total Delay by Number of Trips

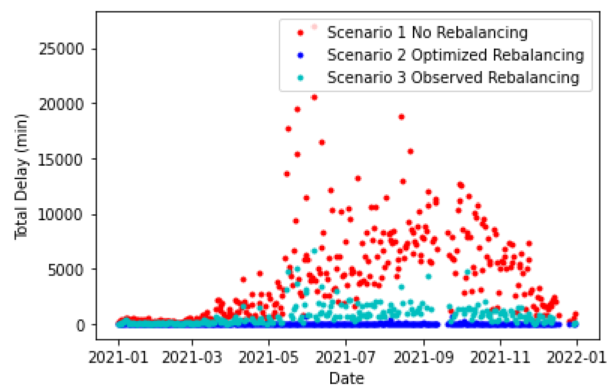

Figure 4: Plot of Total Delay by Date

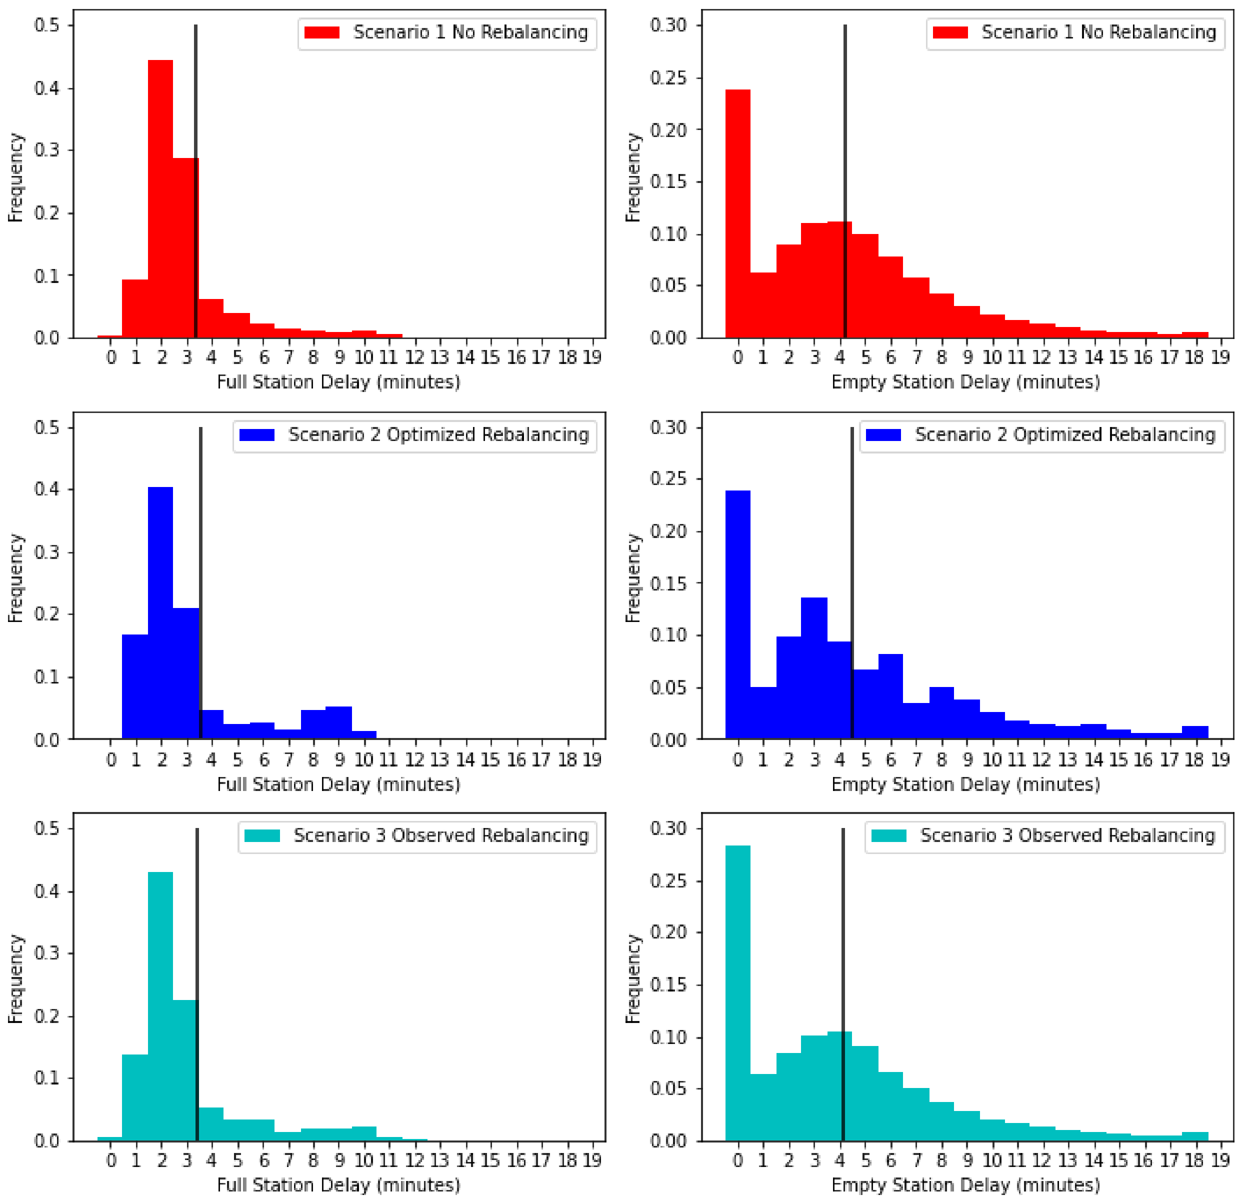

Figure 5: Distribution of Full and Empty Station Delays by Scenario

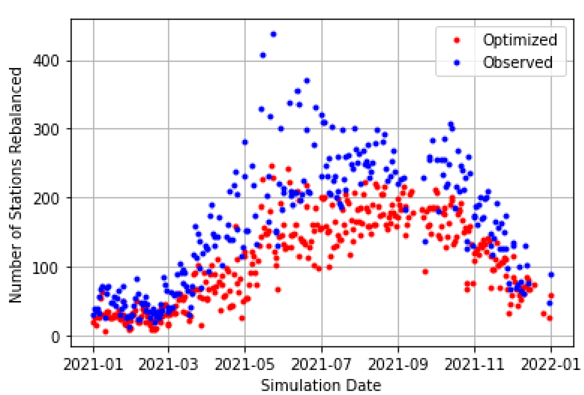

Figure 6: Plot of Station Rebalances per Day

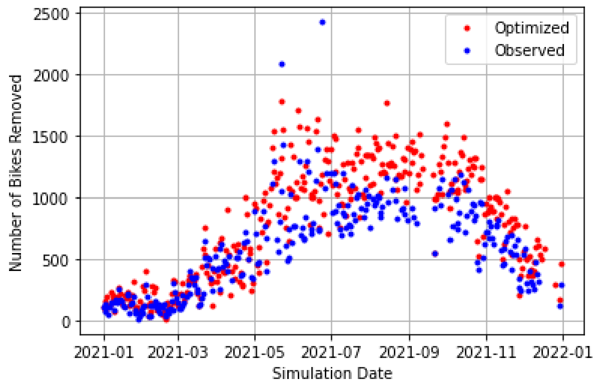

Figure 7: Plot of Number of Bikes Removed from Full Stations per Day

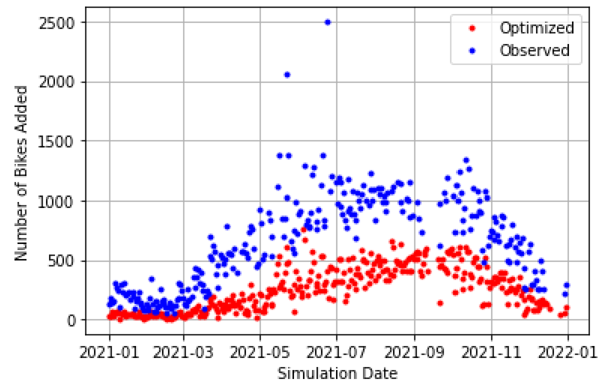

Figure 8: Plot of Number of Bikes Added to Empty Stations per Day

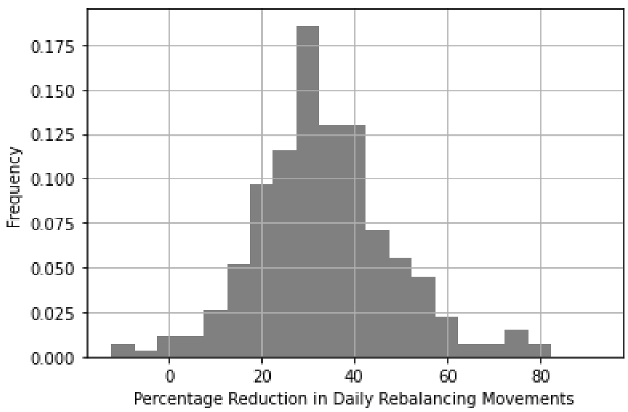

Figure 9: Distribution of Reduction in Daily Rebalancing Movements

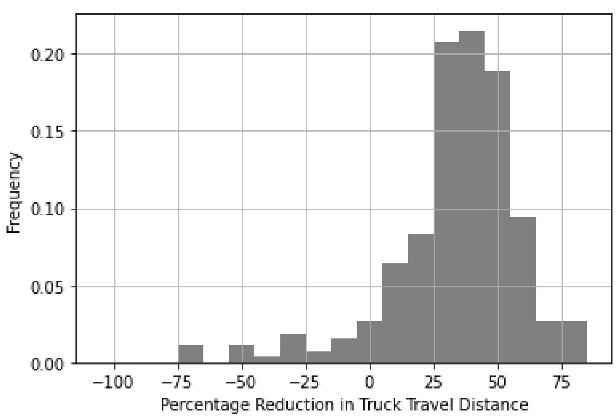

Figure 10: Distribution of the Reduction in Truck Travel Distance for Rebalancing