A Decision Support Framework for Grid-Aware Electric Bus Charge Scheduling

Geoffrey Pettet*†, Malini Ghosal*, Shant Mahserejian*, Sarah Davis*,
Siddharth Sridhar*, Abhishek Dubey†, and Michael Kintner-Meyer*
*Pacific Northwest National Laboratory, Richland, WA, USA
†Vanderbilt University, Nashville, TN, USA
Email: {geoffrey.a.pettet, abhishek.dubey}@vanderbilt.edu,
{malini.ghosal, shant.mahserejian, sarah.davis, siddharth.sridhar, Michael.Kintner-Meyer}@pnnl.gov

***Abstract*—While there are many advantages to electric public transit vehicles, they also pose new challenges for fleet operators. One key challenge is defining a charge scheduling policy that minimizes operating costs and power grid disruptions while maintaining schedule adherence. An uncoordinated policy could result in buses running out of charge before completing their trip, while a grid agnostic policy might incur higher energy costs or cause adverse impact on the grid's distribution system. We present a grid aware decision theoretic framework for electric bus charge scheduling that accounts for energy price and grid load. The framework co-simulates models for traffic (Simulation of Urban Mobility) and the electric grid (GridLAB-D), which are used by a decision theoretic planner to evaluate charging decisions with regard to their long-term effect on grid reliability and cost. We evaluated the framework on a simulation of Richland, WA's bus and grid network, and found that it could save over $100k per year on operating costs for the city compared to greedy methods.**

***Index Terms*—Spatial-temporal optimization, transportation-grid, electric bus, charge scheduling, decision support system**

I. Introduction

Many municipalities have begun exploring the challenges of converting public transit fleets to electric buses (EVs). EVs increase fleet management complexity since the system now interacts with the power grid. Operators must determine where to build charging stations and when to charge the EVs.

Several works have previously examined EV charge scheduling from the perspective of energy cost optimization. Techniques that have been applied to the domain include genetic programming [1], greedy algorithms [2], linear optimization [3], and solving a Markov Decision Process using policy iteration [4].

While minimizing energy cost is important, it is equally important to consider the strain charging decisions place on the power grid. Bus fast charging may have significant impact on the grid and potentially cause thermal overloading, phase imbalances, and voltage violations [5], [6], [7]. This can lead to insulation breakdown in transformers and result in blackouts in the transformer's service area [8].

To accurately account for the energy needs of buses throughout the day with respect to the grid's condition, an integrated traffic and grid model is needed [9], [10], but is understudied in the context of charge scheduling – a recent paper examines charge scheduling with respect to grid constraints [11], but does not include the grid's health in the cost function or consider the impact of traffic on bus movement, for example. In this paper, we present a principled fleet management policy that incorporates such models. A traffic model provides precise travel time and power use information given different traffic conditions, informing the policy when buses will need to be recharged. A power grid model then quantifies how charging decisions impact the grid at various times and locations, and whether the charging behavior causes equipment thermal overloading or unacceptable system conditions [7], [6].

We present an anytime algorithm which leverages these traffic and grid models to estimate both the long term operational cost and power grid strain of charging decisions. Unlike methods such as reinforcement learning that require extensive offline training, our approach requires no training and does not assume a stationary environment. This is crucial in dynamic city environments where traffic or grid loads can change due to unexpected demand or emergencies.

Contributions: To realize such a charge scheduling policy we (1) construct a traffic simulation of the fleet's operating area to extract expected travel times and energy use for each bus route segment, (2) create a power grid model that captures the effect of a charging policy's demand on the grid, and (3) define a control process that utilizes the above simulations to find an optimal charging policy. We implement this framework on the transit system of Richland, WA as a case study and compare it to a greedy charging approach.

II. Decision Support Framework

Grid aware charge scheduling requires many inter-connected components (Figure 2). A traffic model provides bus State of Charge (SOC) and travel time estimates while a power grid model quantifies the grid impact of charging actions and determines if there are infrastructure constraint violations. These models are incorporated in a simulation of the transit system which is used by a decision theoretic planner to look ahead and determine an optimal charging policy. We discuss each of these components in detail below.

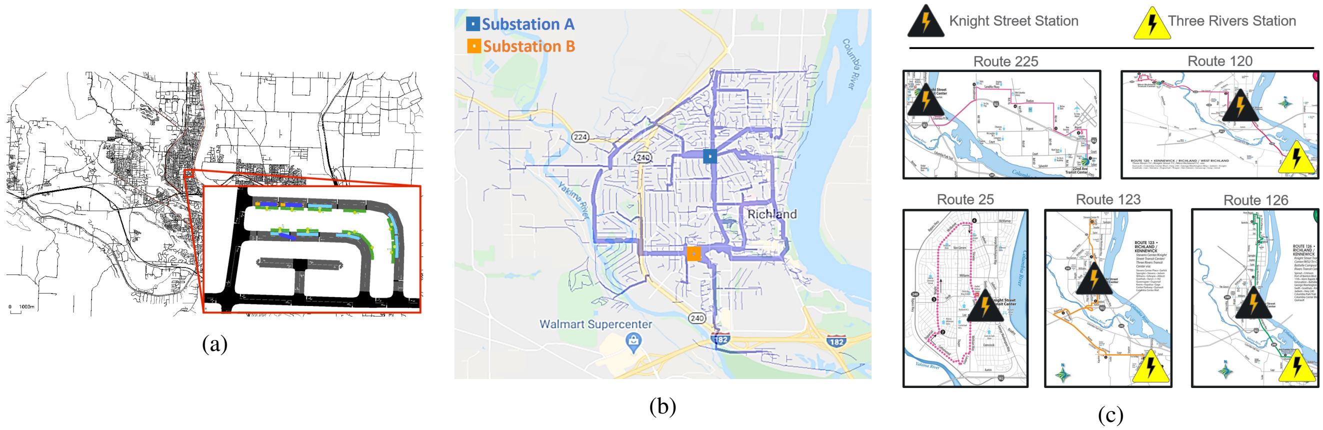

Fig. 1: (a) Street map of Tri-Cities region; the inset screenshot shows buses arriving and parked at the Knight Street Transit Center. (b) GIS overlay of the feeder model. (c) Simulated electric bus routes. Triangles represent the charger locations on each route – black for the Knight Street station, yellow for the Three Rivers station.

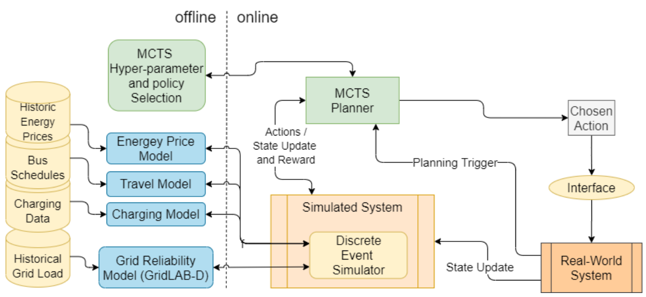

Fig. 2: Grid-aware Decision Support Framework for public transit EV charge scheduling

A. Traffic Model

The traffic model simulates local transit busses and their SOC while following a daily schedule in realistic traffic conditions. The simulation domain for our case study is the Ben Franklin Transit (BFT) service area, which includes Richland, WA, and the encompassing cities of West Richland, Kennewick, and Pasco [12].

The street map data for the considered Tri-Cities region was obtained from © OpenStreetMap (OSM) [13], and included metadata for traffic light programs, public transit routes, and bus stop locations. The street map and zoomed in region of the Knight Street Station transit hub is illustrated in Figure 1a.

The OSM map data was loaded into SUMO, an open source, agent-based, microscopic and continuous simulation package that can handle large networks [14]. The Traffic Control Interface (TraCI) allowed for controlling the SUMO simulation to extract detailed information at every time step in Python. Vehicles in SUMO follow common driving rules, including interactions with other vehicles such as changing lanes and maintaining a minimum space between vehicles [15]. The built in "ElectricBus" vehicle type was used to model the buses and their electricity use, and allows for custom bus attributes such as the attributes listed in [15].

The electric bus flows were generated by repeating routes. Simulation output was collected for the electric buses including travel time between stops and and SOC at each stop.

B. Power Grid Model

The power grid representing the 12.47kV distribution feeder network in Richland, WA is modeled using GridLAB-D [16]. The substation and feeder layouts, equipment ratings, and historical hourly load for every customer was attained from the local utility. In this study scenario, two different distribution substations, each with 12 radial feeders, supply power to the majority of loads within the study footprint. Fig. 1b illustrates the location of these two substations and their feeder network stretching over the city of Richland.

As charging stations are connected to the distribution feeder system, various power system conditions will be monitored throughout the day. Each charging station is mapped to the nearest power grid node, and their impact to power grid depends on their location and power consumption relative to the time of day. Figure 4 illustrates the one-line diagram of the two distribution substations and an example placement of a new charging station placed mid-way down the feeder.

To be consistent with industry practice of building the system to ensure reliability under worst case scenarios, a peak day's hourly load profile is used to represent the power grid load without bus charging. The charging load is then added on top of the peak day load, and load flow analysis is performed to monitor nodal voltage deviations, phase imbalances, line losses, and the apparent power drawn from the feeder head to analyze equipment thermal loading. These measurements are:

- Nodal voltage deviation of the phases $\phi \in \{a, b, c\}$,

$$\Delta v_{i,\phi} = \frac{v_{i,\phi} - v_{base}}{v_{base}} \tag{1}$$

- Imbalance factor [5] of the circuit after charging at node i at time t approximated by

$$\mathcal{I}_i = \frac{v_2}{v_1} \approx \sqrt{\frac{1 - \sqrt{3 - 6\alpha}}{1 - \sqrt{3 + 6\alpha}}}, \tag{2}$$

where,

$$\alpha = \frac{v_{ab}^4 + v_{bc}^4 + v_{ca}^4}{(v_{ab}^2 + v_{bc}^2 + v_{ca}^2)^2}, \tag{3}$$

and v_1 and v_2 are the positive and negative sequence voltage,

- Total line losses L in underground cables (L^{ug}) and overhead lines (L^{oh}) after charging at node i,

$$L_i = L_i^{ug} + L_i^{oh} \tag{4}$$

- Apparent power drawn from the feeder f head corresponding to the node i where the charger is placed at time t,

$$\mathbb{S}_f = \sum_{\phi} V_{f,\phi} I_{f,\phi}, \quad \forall \phi \tag{5}$$

where, the complex voltage and current at the feeder f are denoted by V_f and I_f.

The constraints which should not be violated are:

$$\left.\begin{aligned} |\Delta v_{i,\phi}| \leq \Delta v_{max},\ \mathcal{I}_i \leq \mathcal{I}_{max}, \\ L_i \leq L_{max},\ \mathbb{S}_f \leq \mathbb{S}_{max}. \end{aligned}\right\} \tag{6}$$

where, the suffix max denotes the limit of the measurements. Based on the above measurands in (2)-(5), a novel grid score metric of charging at node i at time t is defined using the following terms,

The grid score g_i is given by,

$$g_i = \begin{cases} 0, \text{ if any of the inequalities in (6) is violated,} \\ 1 - \dfrac{1}{\sum_n w_n}\Bigg[w_1 \displaystyle\sum_{j \in \phi} \frac{1}{3} \frac{|\Delta v_{i,j}|}{\Delta v_{max}} + w_2 \frac{\mathcal{I}_i}{\mathcal{I}_{max}} + w_3 \frac{L_i}{L_{max}} \\ + w_4 \dfrac{\mathbb{S}_f}{\mathbb{S}_{max}} \Bigg], \text{ otherwise.} \end{cases} \tag{7}$$

where, w_n, $\forall n \in \{1, 2, 3, 4\}$'s are the various weights associated with the four additive terms in (7) (normalized voltage violation, imbalance factor, loss and apparent power drawn at feeder head). These weights are introduced so that the planner can choose to prioritize each contributing factor in the metric differently. Uniform contribution from the contributing factors would require each weight to be equal to 1. The joint grid score of multiple chargers' charging impact at a particular time can be derived using an extension of the expression in (7). The time variable t is dropped for simplicity in the above derivation. An example of grid related inputs used are shown in Fig. 3. The individual charging impact is given in Fig 3a, whereas Fig. 3b shows the plot of the assumed time of use (TOU) price of the electricity throughout the day.

C. Decision Theoretic Planner

The overall system goal is to find an electric bus charging policy that minimizes both the impact on the power grid and operating costs. We begin with several assumptions. First, we assume that bus routes are set in advance, that each bus is assigned to a particular route, and that there is a travel model

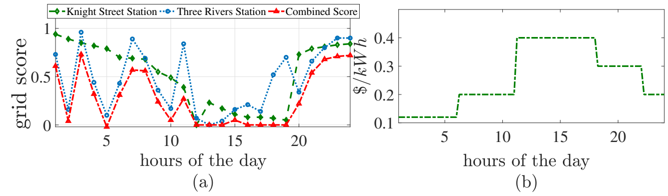

Fig. 3: Grid related input. (a) Individual and combined grid score of the chargers in different hours of the day. (b) The Time of Use price of Electricity.

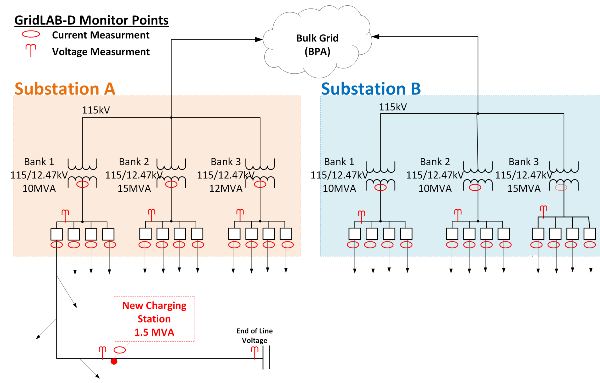

Fig. 4: Richland distribution substations diagram

which describes each bus's travel time and battery discharge throughout their routes (Section II-A). Next, we assume there is a set of pre-defined chargers C placed on the routes. We assume that we are given a model that captures how different charging actions effect the health of the power grid (Section II-B). Last, we assume we have access to a time of use energy price model.

1) Markov Decision Process: We model the bus charging problem as a Markov Decision Process (MDP), which is a model commonly used to describe state and control dynamics for systems with intrinsic uncertainty[17], [18]. MDPs are described by the tuple $(S, A, P(s, a), \rho(s, a))$ where S is a finite state space, A is a set of actions, $P(s, a)$ is the state transition function for taking action a in state s, and $\rho(s, a)$ is the reward function for taking a at s.

States: A state captures environmental information which is needed for decision making, including each bus's SOC and position as well as energy pricing. Our model is limited to states which are relevant to decision making – when buses arrive at and leave chargers. Formally, a state at time t is represented by s^t and consists of a tuple (B^t, ϵ), where ϵ^t is the current time of use energy pricing and B^t is SOC and position information about the set of buses B at time t.

Actions: Actions in our model correspond to assigning a charger $c \in C$ to charge an available bus. A bus $b_i \in B$ is available to charge at c at time t if it is in the set of buses located at c at time t, $\gamma(c, t)$. A valid action at a time t is represented as $a^t = \{c- > b_i | c \in C\}$ where $b_i \in \gamma(c, t)$. $c- > b_i$ represents assigning bus b_i to charge at charger c.

Transitions: State transitions depend on the travel model and charging actions. A bus b_i's location will update based on its position along its assigned route $POS(b_i)$ and the current traffic. We assume that buses do not deviate from their schedule to charge, and spend the same amount of time at the charger weather they charge or not. b_i's SOC change depends on both the travel model as well as charging actions.

Rewards: Our reward function captures an action's impact on both the power grid and operational (i.e. energy) costs:

$$\rho(s,a) = -\bar{\epsilon}_a + \beta \bar{g}(s,a) + \psi n_f(s) \tag{8}$$

where $\bar{\epsilon}_a$ is the total energy cost for taking action a, $\bar{g}(s,a)$ is the total impact to the power grid of taking a at state s (which is mapped to the joint grid score $g(i,h)$ in Section II-B), and β is a hyper-parameter that determines the tradeoff between the two. The last term is a penalty given anytime a bus runs out of charge – ψ is a hyper-parameter, and $n_f(s)$ is the number of 'empty' buses in state s.

2) Solution Approach: When choosing an approach to solve the above MDP, there are two required properties. First, the approach needs to be *adaptive* to unexpected changes in the environment such as equipment failure. Second, it must be capable of handling uncertainty in the environment, including uncertainty in travel times or power grid demand. While the current model does not include such uncertainty, incorporating it is a future goal.

With these requirements in mind, we solve the MDP using Monte Carlo Tree Search (MCTS), a simulation based search algorithm that evaluates actions by sampling from a large number of possible scenarios. The evaluations are stored in a search tree, which is used to explore promising actions. Unlike approaches like reinforcement learning that require offline training, MCTS performs its computation online by sampling from underlying simulations, making it flexible to changes in the environment. There is also substantial research on handling uncertainty with MCTS using techniques such as sparse sampling [19] and information set theory [20]. These properties make MCTS a good choice to solve our MDP.

When implementing MCTS, there are a few domain specific considerations: the Tree Policy and the Default Policy. The Tree Policy governs how the algorithm explores the search tree. We use the standard Upper Confidence bounds applied to Trees (UCT) algorithm [21], which is a principled approach that balances exploiting the most promising actions with exploring other actions. The Default Policy estimates the value of a new node by quickly simulating to a terminal node. The simplest default policy is uniform random action selection, but domain specific information can be incorporated to make these estimates more accurate. Buses with lower SOC's are more likely to be charged at any given moment, therefore our default policy chooses to charge each bus with probability inversely proportional to their SOC.

III. Performance

A. Experimental Design

To evaluate the framework, we examine the Tri-Cities area in Washington, USA. This mid-sized metropolitan area's transit system services the cities of Richland, Pasco, and Kennewick WA. Of the transit system's 18 bus routes, we selected 5 to simulate as EVs, which combined have 14 buses assigned to them on a typical day. We simulated charging stations at two transit hubs – the Knight Street station, which all 5 bus routes pass through, and the Three Rivers station, which 3 routes pass through. This setup simulates one main charging hub, with one secondary hub for a subset of routes, and is shown in Fig. 1c. Our experimental runs are for one day of operation lasting from 6am to 10pm, and assume that each bus starts with batteries at half maximum capacity from overnight charging. To reduce noise, we run 10 experiments for each hyper parameter combination and average their scores.

TABLE I: Experimental Parameters

Hyper-Parameter	**Value(s)**
MCTS Iteration Limit	3000
Look ahead time horizon	3.5 hours
UCT exploit / explore tradeoff	3.5
Bus failure reward penalty ψ	-500
Reward tradeoff parameter β	$\{1,2,3,4,5\}$
Battery capacity	150 KWh
Charging rate	300 KW

For this case study, initial hyper-parameter values and environmental constants were selected from experience and are shown in Table I. We focus on the effect of one key parameter: the reward tradeoff β. It controls the balance between minimizing the system's energy cost with minimizing the system's impact on the power grid, as explained in Section II-C. The other hyper-parameters are kept constant.

To understand the efficacy of our framework, we compare it to a greedy bus charging policy which charges any bus when it stops at a charger if its SOC is under a set threshold. If there are multiple buses that could be charged, the bus with the lowest SOC is chosen. The threshold ensures that buses are only charged when needed. For our experiments we chose a threshold of 41kWh, as this was the lowest threshold that did not lead to bus failure.

B. Results and Discussion

Results are shown in Figs. 5a and 5b. Fig. 5a plots the cumulative energy cost to charge the buses for the 1 day scenario with different values for β, and compares them to the greedy baseline policy. Figure 5b does the same but for the cumulative power grid impact metric. Our first observation is that in all cases our framework outperforms the baseline greedy policy. For energy costs, lower values are better and indicate that less money is needed to charge the buses. The maximum cost using our framework of \$860 is \$50 lower than the greedy policy's \$910. For the power grid metric, a higher value is better and indicates that the charging decisions had a more favorable impact on the grid. Here our lowest power grid score of 376 was better than the greedy policy's score of 362. These values represent a significant savings when scaled to an entire transit system. For example, scaling our results to the full 75 buses in Richland's transit system could save over \$100k per year without considering the avoided cost for grid upgrades or peak demand charges.

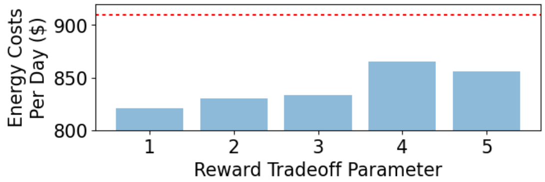

(a)

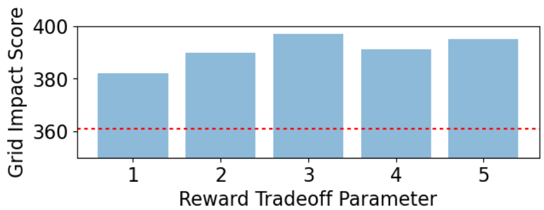

(b)

Fig. 5: Reward tradeoff parameter β's effect on (a) the energy cost (a lower score is better) and (b) the cumulative grid impact (a higher score is better) to run the transit system per day. Red dashed line represents the greedy approach for comparison.

Our second observation is the effect of the reward tradeoff parameter β. Generally, the higher the value of β the more emphasis is given to the power grid impact metric as opposed to the energy cost. This is reflected in our results, as $\beta = 1$ has the lowest grid impact score. As β is increased, the framework increasingly sacrifices energy costs to achieve better grid impacts. The takeaway is that our framework is flexible to the needs of different operators – if a city has very low tolerances for impacts to their power grid but can spare extra operating costs, they can use higher values for β. On the other hand cities with tighter budgets can decrease β to save operating costs at the cost of increasing the stress on their power grid.

IV. Conclusion

When managing an electric bus transit fleet, it is crucial that charging policies take the power grid into consideration. We argue that to realize such a policy requires a decision support framework which incorporates both traffic and power grid models. We discuss how these models are used by a decision theoretic planner that evaluates possible charging schedules with regard to their operational costs and impact on the grid. We implement said framework on mid-sized transit network and found that our approach improves both costs and grid impact compared to a greedy scheduling policy. The positive outcome of this case study motivates future extensions in this domain. For example, our system assumes a conservative worst case scenario of peak historic demand in the grid model – this can be extended to a probabilistic demand model that allows more flexibility in decision making.

Acknowledgement

This research was supported by the Laboratory Directed Research and Development (LDRD) at Pacific Northwest National Laboratory (PNNL), supported by the U.S. Department of Energy under Contract DE-AC05-76RL01830. We would also like to thank the Richland Energy Services (Engineering Department) for their kind cooperation.

References

[1] P. L. Durango-Cohen and E. C. McKenzie, “Trading off costs, environmental impact, and levels of service in the optimal design of transit bus fleets,” *Transportation research procedia*, vol. 23, pp. 1025–1037, 2017.
[2] T. Paul and H. Yamada, “Operation and charging scheduling of electric buses in a city bus route network,” in *17th International IEEE Conference on Intelligent Transportation Systems (ITSC)*. IEEE, 2014, pp. 2780–2786.
[3] S. P. Nageshrao, J. Jacob, and S. Wilkins, “Charging cost optimization for ev buses using neural network based energy predictor,” *IFAC-PapersOnLine*, vol. 50, no. 1, pp. 5947–5952, 2017.
[4] G. Wang, X. Xie, F. Zhang, Y. Liu, and D. Zhang, “bCharge: Data-driven real-time charging scheduling for large-scale electric bus fleets,” in *2018 IEEE Real-Time Systems Symposium (RTSS)*, pp. 45–55, ISSN: 2576-3172.
[5] A. Ul-Haq, C. Cecati, K. Strunz, and E. Abbasi, “Impact of electric vehicle charging on voltage unbalance in an urban distribution network,” *Intelligent Industrial Systems*, vol. 1, no. 1, pp. 51–60, 2015.
[6] S. Deb, K. Tammi, K. Kalita, and P. Mahanta, “Impact of electric vehicle charging station load on distribution network,” *Energies*, vol. 11, no. 1, p. 178, 2018.
[7] M. Kintner-Meyer, S. Davis, S. Sridhar, D. Bhatnagar, S. Mahserejian, and M. Ghosal, “Electric vehicles at scale–phase I analysis: High EV adoption impacts on the western US power grid,” Tech. Rep., 2020.
[8] R. Hermans, M. Almassalkhi, and I. Hiskens, “Incentive-based coordinated charging control of plug-in electric vehicles at the distribution-transformer level,” in *2012 American Control Conference (ACC)*. IEEE, 2012, pp. 264–269.
[9] Y. Lin, K. Zhang, Z.-J. M. Shen, B. Ye, and L. Miao, “Multistage large-scale charging station planning for electric buses considering transportation network and power grid,” *Transportation Research Part C: Emerging Technologies*, vol. 107, pp. 423–443, 2019.
[10] M. Gallet, T. Massier, and D. Zehe, “Developing a large-scale microscopic model of electric public bus operation and charging,” in *2019 IEEE Vehicle Power and Propulsion Conference (VPPC)*. IEEE, 2019, pp. 1–5.
[11] A. Bagherinezhad, A. D. Palomino, B. Li, and M. Parvania, “Spatio-temporal electric bus charging optimization with transit network constraints,” *IEEE Transactions on Industry Applications*, 2020.
[12] BFT Homepage. Ben Franklin Transit Services and Route Schedules. [Online]. Available: https://www.bft.org/
[13] © OpenStreetMap Contributors. Extracted Tri-Cities street map and metadata. [Online]. Available: https://www.openstreetmap.org
[14] P. A. Lopez, M. Behrisch, L. Bieker-Walz, J. Erdmann, Y.-P. Flötteröd, R. Hilbrich, L. Lücken, J. Rummel, P. Wagner, and E. WieBner, “Microscopic traffic simulation using SUMO,” in *2018 21st International Conference on Intelligent Transportation Systems (ITSC)*. IEEE, 2018, pp. 2575–2582.
[15] SUMO User Documentation. Updated: 2020-07-08. [Online]. Available: https://sumo.dlr.de/docs/index.html
[16] “GridLAB-D Simulation Software.” [Online]. Available: https://www.gridlabd.org/
[17] R. A. Howard, “Dynamic programming and markov processes.” 1960.
[18] D. J. White, “A survey of applications of markov decision processes,” *Journal of the Operational Research Society*, vol. 44, no. 11, pp. 1073–1096, 1993.
[19] R. Bjarnason, A. Fern, and P. Tadepalli, “Lower bounding klondike solitaire with monte-carlo planning,” in *Nineteenth International Conference on Automated Planning and Scheduling*. Citeseer, 2009.
[20] D. Whitehouse, E. J. Powley, and P. I. Cowling, “Determinization and information set monte carlo tree search for the card game dou di zhu,” in *2011 IEEE Conference on Computational Intelligence and Games (CIG'11)*. IEEE, 2011, pp. 87–94.
[21] L. Kocsis and C. Szepesvári, “Bandit based monte-carlo planning,” in *European conference on machine learning*. Springer, 2006, pp. 282–293.