



Intermodal Hub Project

Final Report

April 18, 2022

Intermodal Hub Project

Table of Contents

Executive Summary	5
Project Schedule.....	5
1 Task 1: Planning and Site Evaluation	6
1.1 Multi modal charging analysis of power levels and vehicle types	6
1.2 Distribution capacity/needs/impact analysis.....	8
1.3 City and suburban level planning of grid and transportation charging integration	11
1.4 Confirm study participants in addition to UTA (e.g., fleet, including delivery and ride hailing participant vehicles).....	12
2 Task 2: Algorithm Development and System Simulation Analysis	12
2.1 Design initial intelligent prediction algorithms and demand response concepts.....	13
2.2 Develop system simulation models for charging network and agent-based vehicle response.....	16
2.3 Collect data for input to algorithm development and as machine learning training datasets	27
3 Task 3: Software/Hardware Development and Testing	29
3.1 Specify, bid, and procure system hardware.....	29
3.2 Write code and program algorithms on servers to include energy/load balancing and management.....	32
3.3 Evaluate hardware system (with integrated software) at the USU EVR.....	36
3.4 Anticipate needs for and develop cyber security management.....	45
4 Task 4: Piloting and Field Evaluation.....	45
4.1 Accessing and exercising application programming interfaces	46
4.2 Software Architecture for Simulation and Real-Time Control	47
4.3 Sliding-Window Planner for Real-Time Control of Vehicle Charging.....	48
4.4 Field Evaluation	55
5 Cost-benefit analysis	60
5.1 Cost-benefit analysis	60
6 Assessment of the potential for future deployment.....	64

Table of Figures

Figure 1 - Pantograph Overhead Charger	7
Figure 2 – TPSS Electrical Building.....	8
Figure 3 – Diagram of TPSS Power to TRAX.....	12
Figure 4 - Electricity bill breakdown from Nov. 28, 2018 to Dec. 28, 2018	14
Figure 5 - State of charge before and after the DR	15
Figure 6 - State of charge before and after the DR for Bus A	15
Figure 7 – DR Consumption Charge Example Curve.....	19
Figure 8 – Plot of $b(y(t))$	21
Figure 9 – Energy Delivered over Battery Charge.....	24
Figure 10 – Percent Agreement over Episode Number	27

Figure 11 – Proposed Power Monitor Installations.....	29
Figure 12 – Network Diagram	30
Figure 13 – Proposed Data Collection System	30
Figure 14 – Communication Architecture	31
Figure 15 - Comparison of Standard (left) and Smart (right) Peak Management.....	36
Figure 16 - Illustrates Smart Peak Management Combined with Smart Charging.....	36
Figure 17– Boot Notification Response.....	37
Figure 18 – Status Notification Response.....	37
Figure 19 – Get Configuration Response.....	38
Figure 20 – Meter Value Before Charge Limitation.....	39
Figure 21 – Meter Values After Charge Limitation.....	39
Figure 22 – EVR Testbed Diagram.....	40
Figure 23 – A snapshot of new-flyer data used in the algorithm	41
Figure 24 – Small part of route 2 data collection.....	42
Figure 25 – EVR Load for 5 days.....	43
Figure 26 – EMS result for 02/22/2021	43
Figure 27 - Visualization of routes tracked using UTA APIs.....	45
Figure 28 - Architecture for the Bus simulation. Nodes are depicted as rectangles, pub-interfaces are depicted as solid lines, and service calls are depicted as dashed lines.....	48
Figure 29 - Planning framework consisting of a long horizon and short horizon (receding horizon planner, RHP)	49
Figure 30 - Plot of TPSS load over one day (24 hours). The “spiky” nature of this load is difficult to visualize at this scale.....	52
Figure 31 - TPSS load over one hour (6 – 7 AM). The “spiky” nature of this load is apparent.....	52
Figure 32 - TPSS load over 6 minutes. On this time scale, the intermittent structure of the TPSS load may be appreciated.....	53
Figure 33 - A simple graph of the states and transitions for a single charger type. The bottom row corresponds to the charger at rest while the other rows each correspond to charging a particular bus. Yellow ovals are placed around charging windows.....	54
Figure 34 - The top image shows the charge for each of the 30 buses over a 24 hour period. The bottom left shows the power usage of the original 24 hour plan and the bottom right shows the resulting power usage after execution. In each of the figures, the pink zones correspond to peak charge times.....	55
Figure 35 - Top image shows the planned charge for each bus without consideration of the TPSS. The middle image shows the planned charge for each bus while considering the TPSS. The bottom row shows the planned average power usage with the left being without TPSS.....	57
Figure 36 - Upgraded Red Lion device to TPSS Siemens system. USU connected a cellular device to pull data for algorithm input.....	58
Figure 37 - UTA electric bus at overhead charge station.....	58
Figure 38 - TRAX arriving at central station.....	59
Figure 39 - Connections Established on Secure UTA Charger Devices.....	59
Figure 40 - OCPP server script running at USU.....	60
Figure 41 – Script verification that all chargers are connected to USU OCPP server. T175-IT1-003 and -016 are USU ABB chargers. NAMHVC150 and HVC150 are UTA overhead chargers.....	60
Figure 42 - The three charts above show expected impact of estimated annual revenue from electricity on offsetting annual payback on substation infrastructure investment to the utility. Scenario 1 and 2 correspond to the added electric load from expanded bus char	63
Figure 43 - Example opportunities for future expansion of the Intermodal Hub Project technologies with multi-modal support of rail, bus, and private and fleet vehicles. (a) Expansion throughout Salt Lake City at additional TRAX line stations in support of local urban traffic and (b) future batter-electric train replacement of FrontRunner along the I-15 corridor in support of broader passenger and freight electrification in the state.	66

Table of Tables

Table 1 - Project Schedule 6
Table 2 – Effects of DR 16
Table 3 – Project Data Points..... 28
Table 6 - Load scenarios as defined for the traditional cost-benefit analysis 61

Executive Summary

The Intermodal Hub Project successfully developed and demonstrated a power balance and demand response system for multi modal vehicle charging at sites with high peak power demand. The primary goals achieved were to safely and securely manage peak loads, increase utilization of the distribution system, reduce or defer infrastructure upgrade costs, and reduce operating costs. The project was part of the Sustainable Transportation Energy Plan (STEP), and was sponsored by Rocky Mountain Power (RMP) in partnership with the Utah Transit Authority (UTA) and Utah State University (USU).

The project resulted in detailed evaluation of the UTA Intermodal Hub site and its distribution infrastructure, electric loads, and electric bus charging equipment, development of data collection and real-time communications plans, and development of mathematical models, reward systems, and training algorithms within the reinforcement learning algorithm development process. The developed systems were deployed and evaluated at the USU Electric Vehicle and Roadway (EVR) test site and at the UTA Intermodal Hub Central Station. Results demonstrated the ability to significantly reduce peak demand with coordinated charging loads. A cost-benefit analysis was performed to show the return by providing EVs as a resource to improve utilization of charging and utility infrastructure. The project results show tremendous potential for future deployment to reduce operating costs and defer infrastructure upgrade costs at UTA sites along I-15 and at other RMP customer sites with high peak demand and low utilization that operate near high traffic areas.

The team found that existing EV charging infrastructure hardware, software systems, and cloud services are not out-of-the-box ready or compatible with the real-time communications needed for site level dynamic load management. Significant effort was required to work with the commercial overhead and depot charger units as they were not wired or configured for real-time data access, and their cloud services were not configured for low latency communications or even low-level access and control. Similarly, the commercial e-bus cloud service was intended for a high-level customer interface and had significant limitations for real-time access to bus information. These challenges were addressed by working closely with the charging unit supplier and vehicle OEM to gain low level access through APIs and OCPP server development. Early in the project, discussions were also conducted with a potential public EV charger network partner. However, it was quickly realized that their API systems were not currently compatible with real time user site level access, and control to the underlying charging infrastructure through their cloud service or via OCPP was not possible without significant modification. In all cases, it was clear that the partners intend to provide improved access and control in the future and are interested in leveraging the results of this project to motivate internal decision making on future product updates and releases. Additionally, the team identified challenges in modeling and actively controlling the complex behavior existing and envisioned at the UTA Intermodal Hub site. One challenge is in properly defining and modeling the states and the rewards associated with the intermodal hub program within the framework of the reinforcement learning algorithm development process and hard and soft constraints. These challenges were addressed by breaking the problem down into simplified steps and building upon the most important aspects first.

Project Schedule

A Level 1 schedule for the project can be seen below in Table 1. Each project task (i.e. 1-4) was further divided into additional subtasks. These subtasks are explored in more detail throughout the body of the report.

As mentioned above, despite external and unforeseen pandemic circumstances, COVID-19, the project completed the schedule of tasks as outlined with managed delay. A primary constraint for the project was the inability to meet in-person during Tasks 1 and 2. However this was resolved as feasible through remote

opportunities and support with RMP, UTA, and USU personnel. Additional challenges were recognized in Task 4 with software deployment at UTA, specific to the communication interface with the overhead chargers.

Table 1 - Project Schedule

Project Task	2019		2020				2021			
	Q3	Q4	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4
1. Planning and Site Evaluation										
2. Algorithm Development and System Simulations										
3. Software/Hardware Development and Testing										
4. Piloting and Field Evaluation										

1 Task 1: Planning and Site Evaluation

Within Task 1 additional subtasks were identified. Details around our findings, and the individual subtasks, can be found in the subsections that follow. The objective of this task, Planning and Site Evaluation, was to obtain the level of understanding required to effectively map the problem and lay out the roadmap for our algorithm development to support power balance and deployment of a demand response system at the site. Coordination efforts took place with project partners, RMP and UTA, and data collection site visits occurred.

1.1 Multi modal charging analysis of power levels and vehicle types

Introduction

At the Intermodal Hub site, there are currently two types of electric vehicles, namely electric buses and the TRAX light rail, both of which are operated by the UTA. The electric buses can charge using chargers in the bus depot across from Salt Lake Central Station, or they can use a higher-power overhead charger at the station itself. The TRAX trains on the Blue Line receive power from an overhead contact system (OCS), which is powered by a TRAX Power Substation (TPSS) onsite.

Bus Information

The UTA is in the process of acquiring a fleet of electric buses and has started with three New Flyer Xcelsior XE40 buses. Each bus has a battery capacity of 388 kWh, and uses approximately 1.8-2.2 kWh per mile, which varies due to traffic conditions and outside temperature. In the near future, UTA has plans to acquire more than thirty battery-electric buses for use in Salt Lake County, as well as additional depot and high-power overhead charging capabilities at multiple sites.

Depot Chargers

When not in service, the buses are parked in the UTA Central Garage bus depot across 200

South from Salt Lake Central Station. There are three ABB chargers in the depot that share a single 150 kW ABB charger panel. For the purposes of this project, USU has modeled the existing bus depot as presented. However, UTA is currently in the construction phase of a new bus depot. USU has participated in discussions with UTA and RMP to understand upgrade costs, and how the efforts of this project could better assist with future expansion decisions.

Overhead Chargers

UTA has installed two overhead bus chargers at the Intermodal Hub site for the electric bus fleet. Each is rated to charge the buses at 450 kW using a pantograph on top of the bus, as shown in Figure 1. The overhead chargers are meant to quickly charge the bus while waiting for passengers to embark and disembark the bus when stopped on route.



**Figure 1 - Pantograph
Overhead Charger**

TRAX Power Substation

The TPSS is located at the south end of Salt Lake Central Station, as shown in Figure 2. It connects directly to the Rocky Mountain Power 12.47 kV distribution line on 600 West, and then uses internal transformers and a rectifier to produce 750 V direct current power for the TRAX. The TRAX has regenerative braking, so it returns power to the TPSS as it decelerates. When the TRAX accelerates, it draws approximately 1,855 kW of power. As it decelerates, it returns approximately 875 kW to the TPSS. Between Salt Lake Central Station and Old GreekTown Station, for example, it is estimated that the TRAX consumes about 5.5 kWh of net energy, although the TRAX begins receiving power from the next TPSS by the time it reaches Old GreekTown.



Figure 2 – TPSS Electrical Building

Other Loads

In order to keep the station free of snow in the winter, there is a snowmelt system that uses electricity to heat and melt the snow. However, this system is currently out of service for the foreseeable future. The snowmelt systems are connected to several different meters around the site. Using detailed 15-minute data for 2018 and 2019, it was determined that one snowmelt had a peak power consumption close to 400 kW, which is a significant load that must be considered into a monitoring and control system for the site, if the snowmelt is to be used in the future.

Another significant load associated with the site is the CNG compression and refueling station on the west side of the FrontRunner tracks. The refueling station is backed up by an 1100 kW diesel generator, which implies a significant load.

1.2 Distribution capacity/needs/impact analysis

Introduction

In order to quantify the effects of electric vehicle loads at the Intermodal Hub, it is necessary to have a model of the distribution network between the site and Rocky Mountain Power's Fifth West substation. Commercial distribution network analysis is carried out by Rocky Mountain Power using the CYME software. USU does not currently have a license to the CYME software and has converted the CYME database files shared by RMP into a format that can be used by a similar distribution analysis software called OpenDSS that is free and open source. The conversion was carried out using the National Renewable Energy Laboratory's Distribution Transformation Tool.

Analysis in OpenDSS

USU was given CYME database files pertaining to the distribution network surrounding the Intermodal Hub, including the station and all other distribution lines connected to the Fifth West substation. The Intermodal Hub is near the southern edge of that distribution network, and appears to receive power from both the north and south end of the site. UTA has shared several one-line diagrams of the site which USU has used to determine how the major loads (TRAX, overhead chargers, and snowmelt) are connected to the Rocky Mountain Power distribution network.

The database files from Rocky Mountain Power contain fixed peak load estimates for the utility customers on the distribution network. The load estimates were used in OpenDSS to calculate the typical conditions of the distribution infrastructure, including line voltages, line currents, line losses, and active and reactive power on the network. USU analyzed the existing loading information for the site. The given load data indicates a typical load of around 6 MW on the distribution network, which shows that the TRAX, bus chargers, and other Intermodal Hub loads represent a significant portion of the distribution network's total load.

Overview of Modeling Capabilities in OpenDSS

USU has used the Component Object Model capabilities of OpenDSS to automate modeling and analysis of distribution network behavior using scripts written in the Python programming language. Scripting allows USU to perform many functions necessary for creating a realistic model of the distribution network, including:

- Estimation of non-UTA loads on the network using historical load data provided by RMP
- Performance of timestep simulations of the network over durations ranging from seconds to months and recording model outputs to data files
- Measurement of UTA loading impact on the total losses, reactive power consumption, voltage regulation, load factor, and capacity availability of the distribution network

The methods described above were applied to tasks 2 and 4 of this project, since they can quickly run simulations of various simulated or real loading conditions from real-time or historical datasets and evaluate the effectiveness of the algorithms on reducing power costs for UTA and reducing grid impact or overbuild of infrastructure on the RMP distribution network.

Distribution Network Load Estimation

Due to hourly and seasonal variation of demand for electricity, use of a fixed demand value for each load on the network will not provide a realistic estimate of grid impact for most of a given day or year. To properly measure strain on the grid caused by large loads like EV charging, it is necessary to use timestep data that gives an estimate of power consumption for a given time at any point in the year.

RMP has provided USU with electric current data from January 2018 to December 2019 for each phase of the Fifth West substation which powers the distribution network leading to the Intermodal Hub site. USU converted this data from current measurements with a variable sampling rate to substation power data with a fixed 15-minute sampling rate. When an OpenDSS simulation is conducted using scripting, the program uses the Fifth West substation data to estimate the power demand of the non-UTA loads on the network. This is accomplished by looking at the date and time of the simulation and taking the average of the network power load at that time on similar days.

For example, if the time used for the simulation is Friday, March 5, 2021 at 9:30 AM, the script will get the network load data for 9:30 AM on Monday, March 5, 2018, Friday, March 2, 2018, Tuesday, March 6, 2018, and more nearby days in 2018 and 2019. Saturday and Sunday are excluded since weekday load shapes are different from the weekend, and likewise weekday load shapes are excluded for weekend load estimates. Taking the average of the 9:30 AM power demand, the script calculates a load multiplier to proportionally alter the given fixed value for each non-UTA load on the network.

Load estimation is important because high-power EV charging will pose a significantly higher risk of overloading the grid at times of high demand, like hot summer days, but will pose very little risk of exceeding capacity limits at times with minimal non-UTA demand, like the early morning hours in more temperate seasons. Increased flow of power to non-UTA loads will also increase power losses on the grid for currents to the UTA site.

For estimating the effects of time-of-day loading on losses and voltage regulation on the network, taking the average of similar days and times is likely the better way to estimate network load. However, for estimating load factor and capacity availability, it is likely better to use the maximum of the loading at similar dates and times to measure UTA's impact at the most congested possible conditions.

Timestep Simulations

As mentioned in the previous section, use of fixed values of loads is not sufficient for creating a realistic model of the distribution network across hours and seasons, particularly for the UTA site. For example, the TRAX power supply alone can draw nearly 2 MW of power for a few seconds, before using regenerative braking to return hundreds of kilowatts of power to the grid, and then using nearly zero net power. In order to properly measure the impact of the TRAX for every second it is connected to the distribution network, it is necessary to use a time-varying representation of its load. It is necessary to use time-varying loads for the UTA site and non-UTA loads on the network.

The scripts developed by USU take load data for any number of loads on the network in a comma-separated value file and use this data to perform time-varying simulation. For each time in the file, the script calculates non-UTA loads and sets their values accordingly in OpenDSS, and reads the UTA load data from the file and sets their values. Then, by running the OpenDSS engine, the distribution network is simulated, and the script can return almost any parameter of interest to an output file.

Key Measurements of Grid Impact

The term "grid impact" has many potential definitions which will vary depending on the area of concern desired to be analyzed. USU has developed functions that can measure various metrics that indicate grid conditions and UTA's grid impact. A constraint of the distribution network is the amount of current a given piece of equipment can safely handle. Equipment currently must be sized to carry the maximum possible current on the network, so it is important to know how close the equipment gets to their limit during operation. Scripts developed by USU can get the current in each element after a simulation and identify which elements are closest or possibly exceeding their limits.

Another constraint is to always maintain voltage between 0.95 and 1.05 per-unit under all loading conditions on the network. Using the Python-OpenDSS scripts, all bus voltages on the network can be reported, and they can be analyzed in various ways, by creating graphs of the voltage or by using data analysis methods to determine median voltage, standard deviation, etc. At all points along the distribution network between the Fifth West substation and the Intermodal Hub's major loads, the current rating of the distribution equipment is over 100% higher than the expected peak loading on the system. This indicates that in terms of power flow, the system is significantly oversized compared to the loading, even the significant loads from the TRAX and fast charging for electric buses. However, the distribution network is not built only to ensure lines are sized to safely carry current, it is also important to keep voltage

within +/-5% of nominal. At peak load, simulation shows that voltage declines by 1.8% from the substation to the Intermodal Hub. Loads at the Intermodal Hub will also impact voltage at locations on the other end of the distribution network, since they share several hundred meters of distribution lines from the substation. At the far west end of the distribution network, voltage declines by about 2%. Likewise, the voltage declines by about 3.9% at the northernmost extent of the network. Infrastructure upgrades required for the hardware at the Intermodal Hub were not necessary for ensuring power can flow down their distribution lines to charge EVs, but for ensuring the voltage stability across the network when UTA demand is highest. This includes upgrading the lines out of the substation that are shared by all loads on the network, and ensuring the regulators and capacitor banks can keep the voltage within its proper range. Given the several seconds duration of TRAX arrivals and departures, and the several minutes charging sessions of the bus fast charging, these are not just quick transients that will have a very short impact on voltage profiles, they will impact voltage as long as they are being used. Even with shifting charging to occur only when the TPSS is not also active, the systems will still result in slightly depressed voltages across the network (about 0.2%).

Active and reactive power consumption by power delivery elements are also of interest. By reducing peak power demands while keeping total energy consumption the same, it may be possible to reduce energy loss on the distribution network and reduce the need for reactive power compensation. Active and reactive power loss is an output of the simulation that is exported from the OpenDSS simulations and the total kWh or kVAr can be calculated, allowing for energy loss to be compared for different loading scenarios.

A measure of the utilization of the distribution network is load factor, defined as the ratio of average power demand to peak power demand. This can also be calculated using the Python-OpenDSS timestep simulation scripts. Load factor can show how well the existing infrastructure and can identify which loading scenarios make the most effective use of the infrastructure and reduce need for infrastructure upgrades.

1.3 City and suburban level planning of grid and transportation charging integration

TRAX

Like the overhead charging of the buses, the TRAX relies on the overhead contact system for its traction power. There are several TPSS units along each TRAX line. The overhead contact wires are sectionalized using section breaks, which give each TPSS its own section of overhead line to power, instead of each TPSS operating in parallel along the Blue Line. USU looked into the location of the section break between Salt Lake Central and Old GreekTown station to determine the power use at the Intermodal Hub site. The section breaks can be bypassed using a disconnect switch, which would allow the TRAX to run using power from a neighboring TPSS instead of the one closest to it. However, if two consecutive TPSS units are not capable of supplying power, the TRAX cannot run on that section of track.

The Intermodal Hub is the northwest terminus of the TRAX Blue Line, passing through downtown Salt Lake before continuing south to its other terminus in Draper. When connected to the TPSS at the Intermodal Hub, power is drawn from the RMP distribution network connected to the Fifth West substation. When the TRAX reaches the next section of line, the nearest TPSS is receiving power from an RMP substation near the Vivint SmartHome Arena.

In short, the TRAX Blue Line receives power from several substations along its route through the Salt Lake metro area, but only draws power from the Intermodal Hub when it is on the section of line nearest the site.

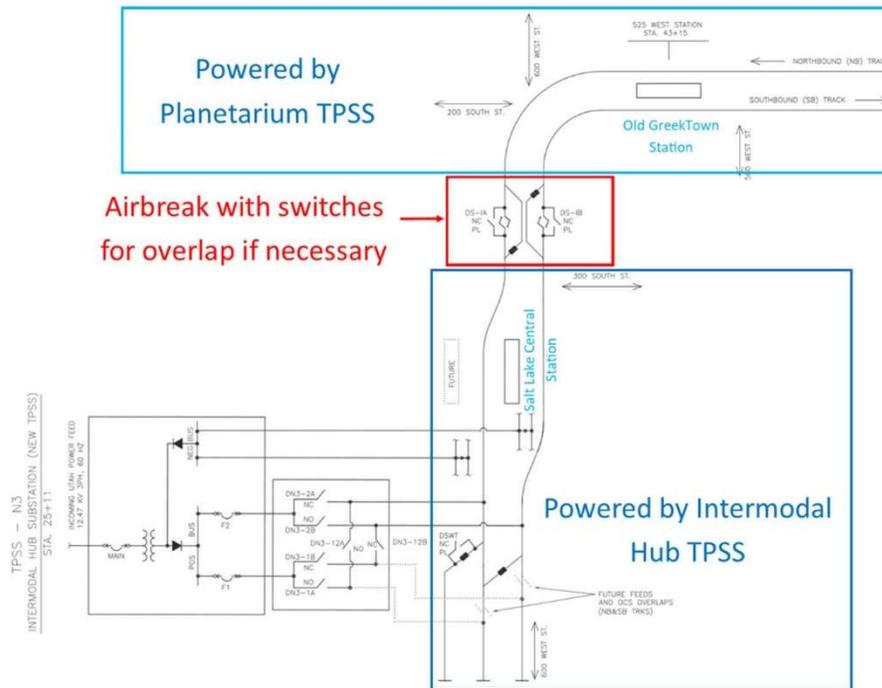


Figure 3 – Diagram of TPSS Power to TRAX

Electric Buses

Currently, the electric buses only charge at the Intermodal Hub site, using the overhead chargers when stopped for passengers, and the depot chargers when the bus is not in service. The buses charge at no other locations on their route, so all energy consumed while driving is energy from chargers at the Intermodal Hub.

1.4 Confirm study participants in addition to UTA (e.g., fleet, including delivery and ride hailing participant vehicles)

A site review for feasibility of EV public access and control has been done, with representatives from USU, UTA and RMP present for a site walk. Initial discussions were had with EV charging equipment vendors (ABB) and third-party EV managers (Greenlots, EV Connect) to understand limitations of current management software and identify requirements for active control through USU developed algorithms.

2 Task 2: Algorithm Development and System Simulation Analysis

Task 2 is decomposed into four subtasks and these subtasks are described in the subsections that follow. Section 2.1 describes the design of an initial intelligent prediction algorithm for demand response. The main conclusions of this section are that (1) optimization is effective at identifying charging schedules that save significant energy and money, and (2) optimal charging may be achieved

by making only minor adjustments to the schedule. Whereas section 2.1 establishes that optimization is worthwhile, the optimization framework applied there lacks the scalability needed for control of real-world operations at the Hub. Section 2.2 addresses scalability by exploring reinforcement learning a branch of machine learning in which a software agent learns actions that maximize expected accumulated rewards. This section presents a model for electric buses and the charging network. Reinforcement learning is briefly described, and the status of software development is given. To accurately portray the physical world, the simulation model requires calibration using data. Section 2.3 describes the data needed and provides a table to summarize the current status of data collection efforts.

2.1 Design initial intelligent prediction algorithms and demand response concepts

Prior algorithms for demand response (DR) in optimal scheduling and charging of vehicles are based on linear programming, mixed integer programming, particle swarm optimization, and genetic algorithms. As the number of vehicles and the types of electric chargers participating in the DR actions increases, the need for an effective and efficient algorithm to find near optimal solutions is imperative.

One of the biggest challenges for optimal scheduling of multi-modal electric demand at the Intermodal Hub is handling the dynamic and massive amount of data that is expected to be collected from the dynamic pricing from the utility, and the stochastic power demands from vehicles (TRAX light rail, e-bus opportunity and depot charging, commuter parking, ride hailing, and fleets). The ability to handle stochastic data and provide accurate prediction makes Reinforcement Learning (RL) suitable for this demand response application in multi modal vehicle charging. RL is discussed further in Section 2.2. Before diving into modeling, simulation, coding efforts, and training for RL, it is worthwhile to explore optimization of charging schedules in a simplified model of the Intermodal Hub. This provides a preview of what may be possible with RL and also provides a baseline for performance comparisons. The remainder of this section explores this optimization approach.

Case Study Introduction

Figure 4 shows a typical monthly bill breakdown for the appropriate Meter on site, with rate Schedule 6A. Around 34% of the cost is from the demand charge, which is the maximum amount of power (kW) drawn for any given time interval (typically 15 minutes) during the billing period, multiplied by the relevant demand charge rate (\$/kW). In fact, the demand charge may comprise up to 50% of a typical utility bill. The highest cost of 44% is from the energy charge, which is the TOU rate multiplied by the energy used in a month. By shifting the bus charging sessions to off-peak hours, savings on the energy bill should be achieved.

To illustrate the benefits of implying the DR in the Intermodal Hub, a 3-bus test case system is proposed in this case study. In this Hub, 3 buses (Bus A, B, and C) with a total of 13 daily charging sessions (6 for Bus A, 4 for Bus B, and 3 for Bus C) are considered.

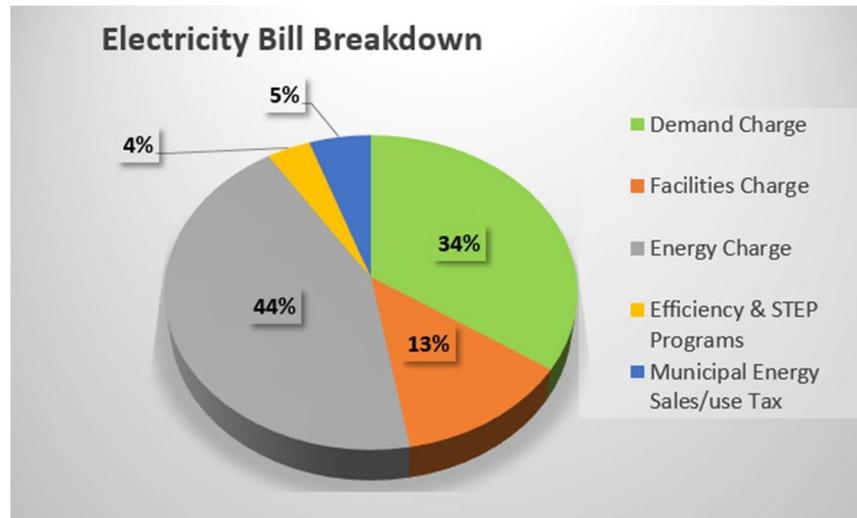


Figure 4 - Electricity bill breakdown from Nov. 28, 2018 to Dec. 28, 2018

Problem Formulation

The three buses run independently and each of them is assigned to an ABB charger in the depot, each of which can charge the bus at 100 kW. The overhead chargers are not included in the model. When not in service or between two scheduled routes, the buses are parked in the garage and ready to be charged, which is represented by a charging window. Each bus has a battery capacity of 388 kWh, and the energy consumption rate is approximately 15-50 kWh per hour, which depends on the traffic conditions and outside temperature. In this model, each bus route randomly generates the energy consumption rate from a uniform distribution over the specified range (15-50 kWh/hour). The DR horizon is a day with 96-time intervals (15-minute). The initial state of charge (SOC) of the bus batteries is set to be 86, 90, and 72%. Using a day-ahead TOU, the DR problem is to find a potentially new start time in each charging window that yields the lowest cost, at the same time, meets the minimum (5%) and maximum (90.4%) SOC constraints. In each charging window, a restricted sliding time window approach is applied to compute all the potential cost and determine the minimum cost that meets the SOC constraint.

The charging window design consists of the choice of the length, defined by the available start and end times. This charging window indicates the shifting boundary. Each available charging window contains a sub-window, representing the charging duration, with the length being equal to the previous (without DR) charging duration. The energy cost is calculated in the sub-window based on the charging rate (kW), charging duration (minutes), and the TOU rate (\$/kWh). The sub-window is then incremented by 0.25 hours, and the process is repeated for the entire time horizon. Each charging section is treated as an independent event. The detailed mathematical definitions of the above-mentioned terms are not presented in this report but are available upon request.

Results

The batteries' daily charging and discharging schedules, with and without the DR, are depicted in Figure 5 and Figure 6 (Bus A only). The discharging periods are represented by the negative slope regions, and the charging periods are represented by the positive slope regions. The flat portion of the curve means the bus is in the garage, either waiting for the next route or to be connected to the depot charger. The model only optimizes the charging start time, all the buses

have fixed service schedules.

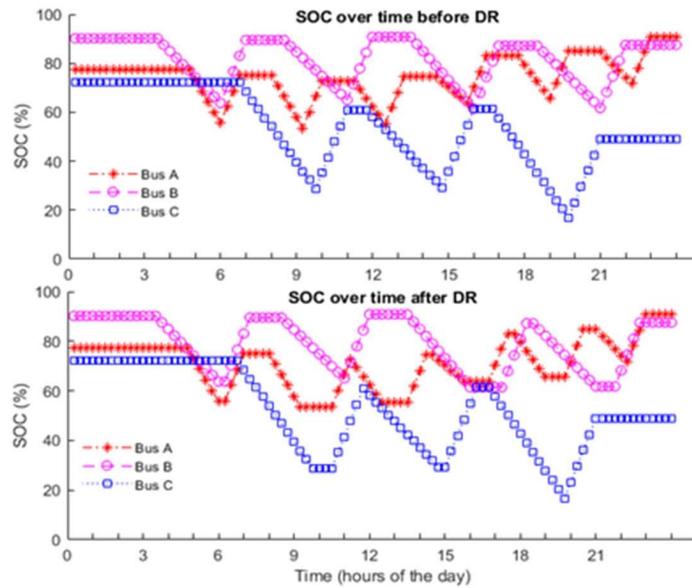


Figure 5 - State of charge before and after the DR

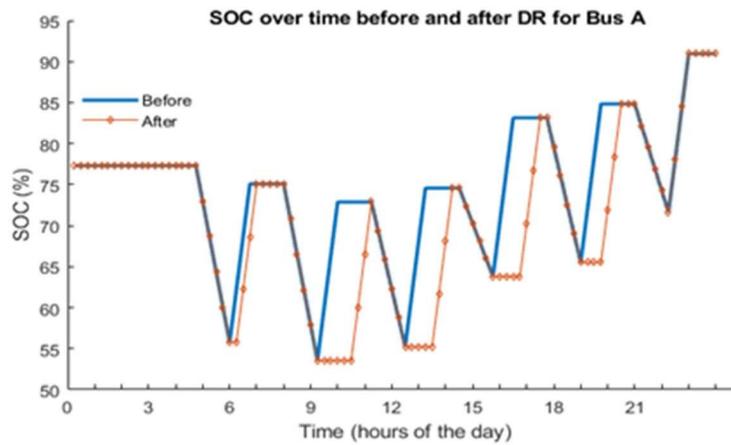


Figure 6 - State of charge before and after the DR for Bus A

To demonstrate the effectiveness of the proposed DR approach, the bus charging schedule is compared with the previous schedule without the DR strategy on daily electricity cost. Simulation results show that the proposed DR algorithm has managed to reduce the total daily cost for Bus A, B, and C by 17.98, 20.87, and 16.87%, respectively as shown in Table 2.

Table 2 – Effects of DR

A	1	06:15-07:00	06:30-07:15			
	2	09:30-10:15	10:45-11:30			
	3	12:45-13:15	13:45-14:30			
	4	16:00-16:45	17:00-17:45			
	5	19:15-19:45	20:00-20:45			
	6	22:30-23:15	22:30-23:15	89.00	73.00	17.98
B	1	06:15-07:00	06:30-07:30			
	2	11:15-12:15	11:15-12:15			
	3	16:15-17:15	17:30-18:30			
	4	21:15-22:15	22:00-23:00	109.00	86.25	20.87
C	1	10:00-11:15	10:45-12:00			
	2	15:00-16:15	15:15-16:30			
	3	20:00-21:15	20:00-21:15	126.00	104.75	16.87

Conclusion

The contribution of the proposed case study is introducing a DR algorithm to minimize the electricity cost without sacrificing the bus service time. Furthermore, the proposed approach is simple in concept and easy to implement. This DR approach has no guarantee of finding a solution that is close to the global optimal. Additionally, this scenario assumes the timing of the bus arrivals and departures are fixed, and the effects of other loads are not considered. In the real world, bus arrivals are inherently uncertain, as are the high power TRAX arrival and departure events against which the bus charging must be balanced to reduce peak power demand. As the types of electricity load and the number of vehicles participating in DR increases, the need for an algorithm to find near-global optimal solutions is imperative. The next step is to apply reinforcement learning on the scheduling of charging.

2.2 Develop system simulation models for charging network and agent-based vehicle response

Introduction

Optimization problems involving demand response and scheduling the charging of electric vehicles have been formulated in a variety of ways. Under simplifying assumptions, the resulting optimization problems may be convex (linear programs) or mixed integer programs, offering optimality guarantees and efficient numerical algorithms. Typically, however, problem formulations are nonconvex and are approached using methods such as particle swarm optimization and genetic algorithms. As the number of vehicles and chargers participating in DR actions increases, the need for effective and efficient algorithms to find near optimal solutions is imperative.

A drawback of optimizing a charging schedule directly as in Section 2.1 is its fragility toward the unexpected. Real-world conditions rarely match the models encoded into optimization constraints and objectives. In a dynamic environment, optimization problems must be re-solved to continuously reflect current conditions. Otherwise performance guarantees are lost, and the solution is only a suggestion to guide actions.

What is needed are methods that scale well to large numbers of vehicles and chargers and are adaptive to changing conditions. Reinforcement learning (RL) has these traits and its performance can improve over time as it learns from experience by interacting with its environment. RL approaches scheduling problems from a new perspective. In RL a software agent observes the state of its environment and learns what actions lead to long-term rewards. RL is patient in the sense that it can allow a short-term penalty if in the long run the action leads to a greater overall return (sum of rewards). RL focuses on choosing optimal actions given arbitrary current conditions. In contrast, optimal scheduling chooses the full sequence of actions assuming everything about the future is fixed and known in advance. With its flexibility, adaptivity, and ability to learn, RL is being used for controlling the Intermodal Hub.

The *agent* in a reinforcement learning problems aims to maximize *return*, which is the expected sum of *rewards*. The reward is the payoff to the agent, and it depends on the *state* of the environment and the *action* that the agent selects. In consequence of the agent's action, the environment transitions into a new so-called *next state*. In mathematical terminology, the environment is referred to as a *Markov decision process* (MDP). A constructive way to introduce RL as it applies to operating the Intermodal Hub is to describe the rewards and returns that the agent seeks to maximize. The independent variables in the reward function define the state of the environment and the actions available to the agent. These variables in turn suggest how to model the environment and simulate the Hub system. Therefore, we begin with a discussion of rewards and returns.

Rewards and Returns for the Intermodal Hub

At a basic level, the quantity to be maximized is money saved. Two types of expenses are initially considered. The first is the money spent on electricity to operate the hub which includes charging bus batteries, melting snow, keeping the lights on at the station, and keeping the TRAX trains running. Operating the hub efficiently means saving as much money as possible on electricity. A second category of expenses are the bus batteries themselves. Batteries wear out over time, but the lifetime of a battery may be prolonged by carefully controlling the battery state of charge (SOC) over its charging and discharging cycles. Charge and discharge rate (fast vs. slow) and state of charge (high or low), and heating are a few of the factors that influence battery health and lifetime. Extending battery life should be rewarded because it lowers overall costs for operating a fleet of electric buses. Other cost saving measures may be included in the reward function. Initially we focus on electricity usage and battery lifetime.

Electricity Costs

The model incorporates electricity costs. Customers are charged based on energy measurements at each electricity meter. Each meter has a schedule associated with it. The schedule describes how charges are assessed by the utility. Components commonly found in schedules include fixed monthly charges, consumption charges, and demand charges. Meters and these schedule components are discussed below.

Power and Energy and Electricity Meters

Let $v(t)$ and $i(t)$ be the total voltage and current drawn at an electricity meter. We adopt the simplified model in which each load connected to the meter experiences the same voltage $v(t)$ but draws its own current $i_m(t)$. The total current is the sum

$$i(t) = \sum_m i_m(t).$$

The total instantaneous power consumed at the meter is

$$p(t) = v(t)i(t) = v(t) \sum_m i_m(t) = \sum_m v(t)i_m(t) = \sum_m p_m(t),$$

where $p_m(t) = v(t)i_m(t)$ is the instantaneous power delivered to the m^{th} load. The total power $p(t)$ is the sum of the powers delivered to the m individual loads.

Power $p(t)$ and energy $e(t)$ are related by

$$p(t) = \frac{d e(t)}{dt}, \quad e(t) = \int_0^t p(\tau) d\tau.$$

The energy consumed at a meter over a month is $e(T)$, where T is the duration of the month.

Consumption Charges

Electric utilities charge customers for energy consumption. Under a simple price schedule, the consumption charge is a rate s_c times the energy used in a month $e(T)$. Define the consumption charge at time t as

$$q_c(t) = s_c e(t) = s_c \int_0^t p(\tau) d\tau.$$

Then the simple consumption charge at month's end is $q_c(T) = s_c e(T)$.

A more complex price schedule may use a high rate s_1 for the first e_1 kWh of energy used and then use a lower rate s_2 for additional kWh. In this case the total consumption charge for the month would be $q_c(T)$, where $q_c(t)$ is computed as

$$q_c(t) = s_1 \max(0, e(t)) + (s_2 - s_1) \max(0, e(t) - e_1).$$

Hereafter this type of consumption charging scheme is referred to as diminishing-rate (DR) consumption charging because the rate decreases as consumption increases. The Figure 7 below shows an example of this consumption charge in which $s_1 = 0.117$ cents/kWh, $s_2 = 0.066$ cents/kWh, and $e_1 = 1500$ kWh.

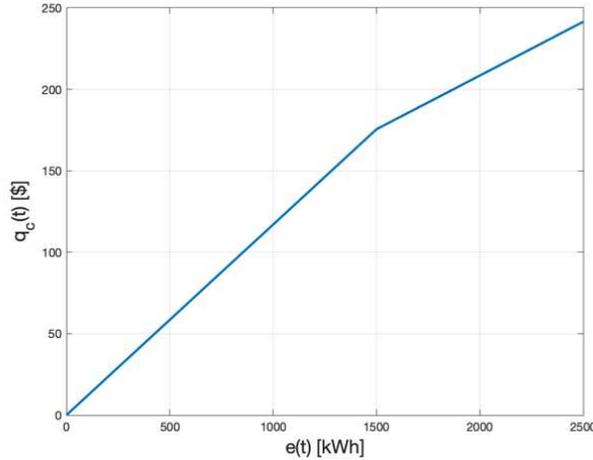


Figure 7 – DR Consumption Charge Example Curve

Under some pricing schedules the rate may vary with clock-time where a lower rate s_{off} is assessed during off-peak hours and a higher rate s_{on} is assessed for on-peak hours. Let $s(t)$ be the rate assessed at time t , then the consumption charge at time t is computed by

$$q_c(t) = \int_0^t s(\tau)p(\tau)d\tau.$$

The total charge for the month is $q_c(T)$. Note that the previous case of a constant rate $s(t) = s_c$ is a special case of this more general expression. Hereafter this charging scheme is referred to as clock-time rate (CTR) consumption charging.

Demand Charges

Utilities may also assess a power demand charge. Power demand at time t is calculated based on a 15 minute (900 second) average power

$$p_{15}(t) = \frac{1}{900} \int_{t-900}^t p(\tau)d\tau.$$

The electric utility must supply enough power to meet the worst case (peak) power demand for the customer

$$p_{\max}(t) = \max_{\tau \in [0,t]} p_{15}(\tau).$$

The demand charge $q_d(t)$ is the peak power demanded by the customer if it exceeds a fixed amount of average power p_{fix}

$$q_d(t) = s_d \max(p_{fix}, p_{\max}(t)),$$

where s_d is the demand rate. The monthly demand charge is $q_d(T)$.

Price Schedules

The price schedules used at the Intermodal Hub include RMP schedules 6A and 23. The TRAX train (one meter) and snowmelt (three meters) utilize schedule 6A. Bus chargers utilize schedule 23. The charge structure for these schedules is summarized below. Fixed customer and facility costs are neglected here. These schedules have different rates during summer (May

through September) and winter (October through April), but for simplicity the description below assumes a fixed rate throughout the year. Schedule 6A has a CTR consumption charge that uses different on and off-peak rates. Schedule 23 consists of a DR consumption charge and a demand charge. The parameters p_{fix}, e_1 and rates $s_d, s_{\text{on}}, s_{\text{off}}, s_1, s_2$ may be looked up in the published price schedules.

Bus Battery Lifetime

Batteries have a natural lifetime and that lifetime can be extended or shortened based on patterns of charging and discharging. The rate (fast vs. slow) of charging and discharging as an effect on battery life. These effects are currently being studied in published literature and will be integrated into the model as appropriate. Very high and very low battery state of charge (SOC) also influences battery lifetime. Let $\gamma(t)$ be the battery SOC at time t . The SOC is related to power consumed at a charger by

$$\gamma(t) = \gamma_0 + \zeta \int_0^t p(\tau) d\tau = \gamma_0 + \zeta e(t),$$

where γ_0 is the SOC at the beginning of the month and ζ is the charging efficiency factor which is typically in the range $\zeta \in [90,95]$ percent.

A battery's lifetime can be shortened if the stored charge $\gamma(t)$ falls close to 0% of capacity or rises near 100% of capacity. Battery lifetime is prolonged if the stored charge lies between γ_1 and γ_2 percent of capacity. Battery harm is assumed to be zero for stored charge in the interval $\gamma(t) \in [\gamma_1, \gamma_2]$ and harm is assumed to increase linearly outside of this interval to a maximum harm of η_1 at $\gamma(t) = 0$ and η_2 at $\gamma(t) = 100$ percent. The η parameters have units of dollars/SOC. We are currently investigating meaningful values for the η parameters based on published literature. The battery harm function may be computed as

$$b(\gamma(t)) = \max\left(\eta_1 \frac{\gamma(t) - \gamma_1}{0 - \gamma_1}, 0, \eta_2 \frac{\gamma(t) - \gamma_2}{100 - \gamma_2}\right).$$

In Figure 8 below shows a plot of $b(\gamma(t))$ for values $\gamma_1 = 50\%$, $\gamma_2 = 80\%$, $\eta_1 = 0.50\$/\text{SOC}$, and $\eta_2 = 0.08\$/\text{SOC}$.

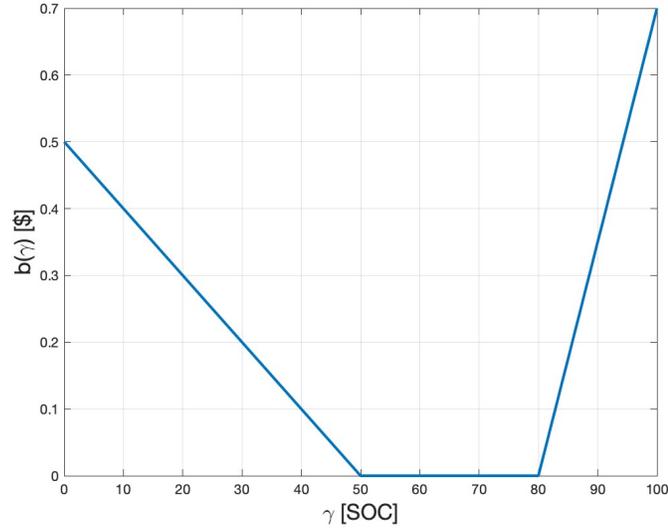


Figure 8 – Plot of $b(\gamma(t))$

Battery harm is cumulative, so the total harm done at time t is given by

$$a(t) = \int_0^t b(\gamma(\tau)) d\tau.$$

The total battery harm for the month is $a(T)$. The agent is incentivized to choose actions that keep the instantaneous SOC between γ_1 and γ_2 . Other effects that affect battery lifetime are currently being studied and this model will be expanded as more is learned.

Return for the Intermodal Hub

Current focus is on the meter supplying energy to the bus chargers. This meter uses RMP schedule 23 which includes a DR (diminishing rate) consumption charge and a demand charge (based on peak 15-minute average power consumption). Combining these charges with the battery harm leads to the total return for the meter over a month

$$\text{return} = -[q_c(T) + q_d(T) + \theta a(T)].$$

The negative sign is introduced to reflect the idea that return is to be maximized and we want to minimize the consumption charge, the demand charge, and battery harm. The $\theta > 0$ parameter weights the relative importance of minimizing battery harm versus the importance of minimizing consumption and demand charges.

Discretization

The software agent interacts with the environment at discrete instants of time $t_0 = 0, t_1, t_2, \dots, t_N = T$. These time instants need not be uniformly spaced samples. Rather they are simply the times at which the agent observes the state of the environment and makes decisions. It makes sense to use uniformly spaced samples during the hours when the hub is in operation and use no samples during night-time hours.

The power, energy, charge, and battery harm relations which are expressed above as integrals in continuous-time may be discretized to fit into the RL framework.

Returns and Reward Functions

In the canonical formulation of reinforcement learning, the return G_n is the cumulative sum of rewards to the terminal state

$$G_n = R_{n+1} + R_{n+2} + \dots + R_N = \sum_{k=n+1}^N R_k = R_{n+1} + G_{n+1}, \quad n = 0, 1, 2, \dots, N,$$

where $G_N = 0$ because the sum is vacuous in that case. The sum of all rewards in an episode (month) is

$$G_0 = R_1 + \dots + R_N = \sum_{k=1}^N R_k.$$

The interesting feature of return G_n is that it accumulates from the terminal state backward in time. In contrast, the consumption charge, demand charge, and battery harm accumulate forward from the beginning of the month. To cast the hub control problem into the RL framework, we equate

$$G_0 = -[q_c(T) + q_d(T) + \theta a(T)],$$

but this leaves us with the problem of finding ways to compute the intermediate rewards R_n and returns G_n . We have solved this problem and will include the mathematical details in a future version of this report. The mathematical model, which includes multiple electric buses and multiple charging lanes at the hub, includes a discrete-time, linear, time-varying state-space model for the evolution of energy, average power, and battery SOC. The consumption and demand charges and battery harm effects are easily computable nonlinear functions of the state variables. The reward and returns are also easily computable from the consumption/demand charges and battery harm.

Environment State Space for the Intermodal Hub

The state space for the RL problem describes the environment. In this case, the environment is the energy consumed at a meter, the peak 15-minute average power, and the battery SOC on all buses in the fleet. We assume that these variables are available to the software agent residing at the hub. The meter readings are local, but information from buses must be communicated back to the hub through wireless data links.

Agent Action Space

The actions available to the agent are commands given to electric buses. As buses approach the hub, the agent may (1) assign a bus to a charging lane, (2) command a bus to stop charging and leave the charging lane, and (3) command the bus to skip charging and move directly to passenger pickup. These decisions affect the energy consumed, the 15-minute average power, and the SOC of bus batteries. Each of these in turn impact consumption and demand charges and battery harm.

Reinforcement Learning

The dimension of the environment state space scales linearly with the number of active buses in the fleet and linearly with the number of charging lanes. The state space variables are continuous in nature. The energy consumed, average power, and SOC are real positive numbers

(as opposed to being integers). The action space, on the other hand, is discrete. There are a fixed number of discrete actions that may be considered at any given time instant. This continuous-state/discrete-action configuration is commonly encountered in RL.

The RL agent learns a policy, which is a function that maps states to actions. Thus, RL is reduced to the problem of function approximation. Many different approximation schemes have been explored, but an approach which is now quite common is to employ deep neural networks to map states to actions and this is the approach we plan to employ in our initial experiments. An optimal policy is one that maximizes the accumulated rewards from all the actions in an episode (month).

Simulation

The software agent learns the policy (neural network) through interaction with a simulated or real environment. We have developed a simulator that evolves the environment state variables described above. Buses traverse routes in the simulation, as they do their batteries discharge. Buses follow the commands of an agent to enter or bypass charging lanes at the hub. Buses in charging lanes have their batteries charged and incur consumption and demand charges at the meter. The mechanics of these things have been worked out. The simulator is written in Python for convenience in interfacing with the machine learning tools.

Scale of Simulation and Modeling

Our initial efforts were focused on RL for a single meter (the bus charging meter). When the RL training, testing, and validation for a single meter has been completed, we plan to scale to hub level control. This testing will continue into Task 3 with the deployment of the algorithm at the EVR. Through the process of scaling, this will allow exploration of new pricing schedules that apply at the hub level and this is where a significant benefit may be realized for both RMP and UTA. The facility to train and validate RL techniques for hub-level control will enable us to explore hypothetical pricing schedules and evaluate the savings for RMP and UTA. Alternative pricing schedules do not affect the underlying model of the environmental state. It only affects the nonlinear functions that transform the state space variables into the quantities involved in reward functions.

Schedules and Model Calibration

Information about the meter schedules applicable at the hub have been received and digested and fit into the model above.

We have received and digested some information relating the time that a bus is in a charging lane and the energy consumed to the change in battery charge. One of these relationships is plotted below in Figure 9 along with a linear regression. This type of data feeds directly into our model and simulation.

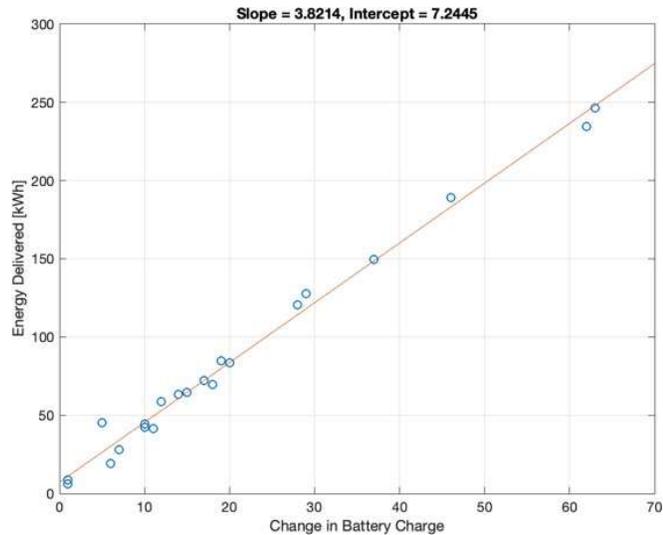


Figure 9 – Energy Delivered over Battery Charge

Adaptation for EVR Implementation

The EVR at USU is a testbed where the learning concepts described above may be tested in the real world as part of Task 3. A battery electric bus simulates traversing an actual route by looping around a quarter-mile track some number of times. This simulates the bus leaving the station and then returning after traveling a predetermined distance and having the expenditure of energy reflected in the reduced state of charge on the bus battery. As at an actual bus depot, the EVR is equipped with a bus battery charger (manufactured by ABB in this case). When the bus returns to the depot, the RL agent decides whether to charge the bus battery or not and if charging is decided, then the agent decides how much or how long to charge, i.e. the agent chooses the incremental increase in battery state of charge.

As described above, the EVR provides the ability to prototype some of the physical realities in operating a real fleet of electric buses. However, it should be understood that the one-bus/one-charger (1B1C) setup at the EVR leaves out some of the realities encountered with an actual fleet. The EVR operates one bus whereas a fleet consists of many busses including a mixture of battery electric busses as well as diesel and CNG busses. The EVR has a single charger whereas a transit system would operate many charging sites with several charging slots at each site. This expanded reality creates new opportunities and challenges such as contention for charging resources and greater need for peak power management and demand response. Performing 1B1C experiments at the EVR is, nevertheless, useful as a proof of concept for reinforcement learning and all of the supporting infrastructure required for data telemetry and charger controls.

The EVR also plays an important role in expanding beyond 1B1C experimentation. To understand how, consider an electrified bus transportation system with N buses (N routes) and M chargers at a central station. We refer to this as the N -bus/ M -charger (NBMC) scenario. This scenario is a close approximation to the full complexity encountered at the Intermodal Hub in SLC. Resource contention, demand response, and peak management are all encountered in the NBMC scenario. So, this is a significant increase in complexity relative to the 1B1C scenario. We can train an RL agent to control bus battery charging in the NBMC

scenario. The capability and performance of the trained agent may be tested using the single bus and charger at the EVR through hardware in the loop simulation. This concept sets up a simulation in which there are N buses. The battery state of charge and distance traveled of $N-1$ of those buses are computed in a computer simulation, and 1 of those buses is a physical bus with real batteries traveling through space in the real world. Similarly, $M-1$ chargers are simulated, and 1 charger exists in the physical world. By feeding data from the real bus and real charger into the simulation, the physical bus and charger may participate in the simulated environment to stress test the trained agent. The software agent does not know the difference between a simulation and the real thing. It simply receives inputs and selects actions that maximize reward. The main difficulty with hardware in the loop simulation is reflecting the contention at the real charger between a real bus and virtual buses and how to simulate energy used when a virtual bus is commanded to connect to the real charger. This non-physical event may be accounted for through bookkeeping in software by providing two energy variables for the real charger: one to reflect the actual energy expended to charge a real bus battery, and one to account for virtual charging. This effort is considered part of Task 3 rather than Task 2, and will be reported on at a later date.

The 1B1C scenario was explored as part of Task 2 efforts. Not only is the 1B1C scenario a simplification of the full NBMC framework, the complexity of the 1B1C scenario offers a closed form solution for optimal charging. In the single bus case, several terms in the reward function are of no use. The consumption charge becomes irrelevant because the total energy required to keep the bus moving is fixed. The demand charge is obviated because the peak power is set by the need to charge a single bus. Two or more buses are needed to make demand charge interesting. This leaves only the battery health to optimize.

Quantize the battery state of charge to 11 levels: 0%, 10%, 20%, ..., 100%. Battery percent discharge d for traversing a route is modeled as a uniform random variable taking value 20%, 30%, or 40%. The battery health function is given by

$$R(s) = \begin{cases} -0.2(s - 54\%)^2 + 1, & s = 10\%, \dots, 100\%, \\ -50, & s = 0\%, \end{cases}$$

where s is the battery state of charge. The battery health is a function of the state, which is taken to be battery state of charge when the bus arrives at the station. The battery health is maximized when the battery state of charge is 54% and decays for larger and smaller states of charge. The large negative reward for 0% state of charge is the penalty assessed to inform the agent that the battery should never be allowed to become fully discharged. The action a available to the agent are the increase in percent state of charge of the battery when the bus arrives at the station. The maximum amount of charge is limited by the length of time the bus is at the station. Therefore, the actions are 0%, 10%, ..., 40%. The next state s' is equal to the previous state s plus the discharge d plus the charge amount a , $s' = s + a - d$. The optimal (greedy) policy for charging the bus may be obtained by repeating until convergence the following value iteration

$$v_{k+1}(s) = \max_{a \in A(s)} \sum_{s', r} p(s', r | s, a) [r + v_k(s)].$$

The resulting optimal greedy policy is given in the table below which shows action a that should be taken as a function of the state.

Optimal Greedy Policy	
State s [%]	Action a [%]
0	40
10	40
20	40
30	40
40	40
50	30
60	20
70	10
80	0
90	0
100	0

This policy illustrates a charging schedule that reflects the uncertainty in the amount of discharge experienced along the route. Notice that computing the optimal policy requires knowing the joint conditional probability distribution $p(s', r|s, a)$ of the next state s' and reward r given knowledge of the current state s and action a . In the simplified learning problem presented here, this probability distribution may be computed. In larger, more complex problems, evaluating $p(s', r|s, a)$ becomes intractable.

Reinforcement learning attempts to learn the optimal policy through simulated experience with the environment rather than relying upon knowledge of $p(s', r|s, a)$. We applied a technique called Double-Q Learning (DQL). This algorithm iterates on the action value function for each state-action-reward-next-state tuple that is observed during the episodes

$$Q_1(s, a) = Q_1(s, a) + \alpha[R(s) + Q_0(s', \arg \min_{a' \in A(s')} Q_1(s', a')) - Q_1(s, a)].$$

The update rule randomly selects to update either Q_1 or Q_0 . The Q_1 update is illustrated above. The randomness breaks a maximization bias present in simple Q-learning and leads to better policies. We ran 1,000 episodes with 1,000 steps per episode, using 5% exploration, and a learning rate of 0.002. Learning curves for this experiment are shown in the Figure 10 below.

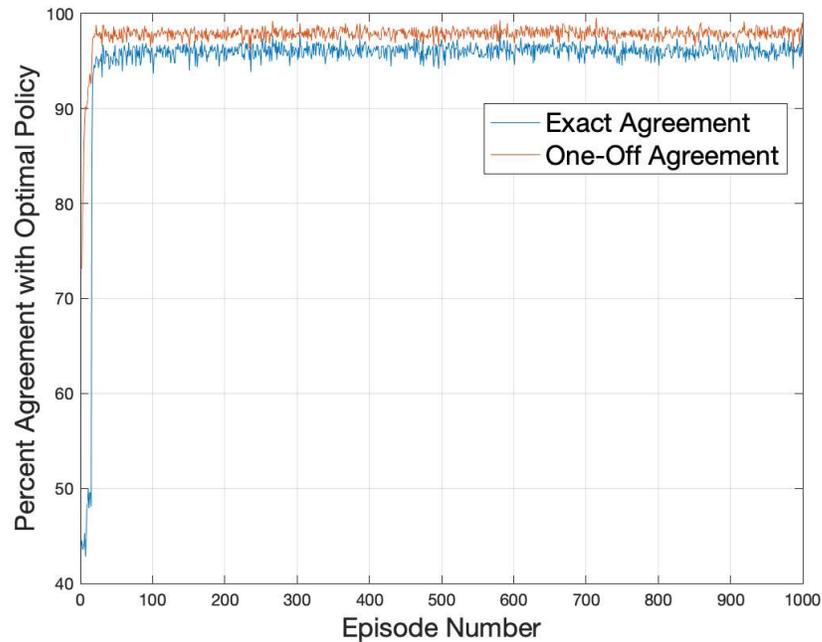


Figure 10 – Percent Agreement over Episode Number

The learning curves plot the agreement between the learned policy with the optimal policy as a function of the number of episodes used in training. The learning curves illustrate that a steady state is reached very quickly, in this case after about 100 episodes. Two curves are shown. The “Exact Agreement” curve shows the agreement between charging actions selected by the optimal policy and the learned policy. The “One-Off Agreement” allows a difference between the decision made by the learned policy to be one charging increment different from the optimal policy. The differences here are to be expected because the policy used during learning uses a 5% exploration policy and the “Exact Agreement” curve tops out at 95%, which is 5% off from the optimal policy. This is to be expected during learning. In real-time operation when exploration is not used, the learned policy would be in exact agreement with the optimal policy. Exploration is needed during learning to allow the agent to consider alternatives that it has not previously considered in order to evaluate and compare them to prior experience.

2.3 Collect data for input to algorithm development and as machine learning training datasets

USU received power data for several types of loads at the site, as well as bus and TRAX schedules for input to the simulation and training algorithms with the goal of integrating real-time data (Task 3 testing).

The TRAX data requires information about the route, and its power consumption. USU used the route schedules posted by UTA as well as estimates of its power consumption to simulate its behavior. TPSS metering is used to monitor when the TRAX is drawing or pushing power to the grid via the TPSS. Additionally, one-line diagrams from UTA have given some information about which section of track is powered by the Intermodal Hub TPSS.

New Flyer Connect reports were used to collect data for the energy consumption of New Flyer

electric buses and to draw estimates of the effects of routes and weather on energy consumption. Initially these reports were generated manually, but USU has worked with New Flyer, and their API, to establish an interface with at a higher sample rate. The reports are for the New Flyer buses only, and with UTA’s forthcoming electric bus procurements considerations will need to be made to obtain required reports from other battery electric bus manufacturers as necessary.

Data from RMP meters around the Intermodal Hub area including, UTA billing information, have given insight into energy usage at the hub, and reports from on-site chargers help with understanding how much power is used and when. The bills can be used to estimate energy use and peak demand using the RMP rate schedules. There is 15-minute timestep meter data for one snowmelt onsite for 2018 and 2019, however the current metering is outlying and not representative of snowmelt conditions, as the snowmelts are out of commission until further notice.

This data and other information such as bus and TRAX route information are fed into the simulation, with the goal of matching real-world scenarios as closely as possible.

Upgrades were made onsite at UTA used to monitor power consumption of the TPSS.

Table 3 – Project Data Points

Required Data	Data Received	Sample Rates	Data Collection Plan
TRAX - Route Info	Schedules from UTA	N/A	Complete
TPSS Power Profile	Estimate of load profile, billing info	N/A, monthly	Metering of TPSS – Siemens system upgrade completed with UTA
E-Bus (New Flyer) - SOC and location	New Flyer Connect reports	Manual	API access and historical/real-time information logged
CNG Compressor Station	Billing data	Monthly	Onsite metering – installation pending contractor bid package. Historical data has been used to create and trend models
Snow Melt	Load data for one unit in 2018 and 2019	15 minutes	Onsite metering – currently not functional with no immediate plans to correct. Historical data has been used to create models
Bus Depot Chargers	ABB charge reports	Manual / Real-time with OCPP integration	OCPP server to pull charging information real-time.
Overhead Charger	New Flyer Connect reports	Manual / Real-time with OCPP integration	OCPP server to pull charging information.
Rate Schedules	RMP Rate Schedule document	N/A	Complete
Infrastructure Upgrade Costs	Conversation with RMP	N/A	Complete
Substation/Distribution Network	CYME files from RMP	N/A	Complete
Other Site Loads	Billing data	Monthly	Complete
Intermodal Hub Operational	RMP Rate Schedule document	N/A	Complete

Information			
UTA Bus Route data	Some available on UTA web site (departure times at various sites)	N/A	UTA API for learned and forecasted route planning

3 Task 3: Software/Hardware Development and Testing

3.1 Specify, bid, and procure system hardware

Part of the system hardware to be considered at the UTA Intermodal Hub site are power monitors. Power monitors are devices which give real time information about the power system to which they are connected. These parameters can vary based on the monitors chosen and the extent of information needed from the power system. Examples of the data to be collected from the UTA site include:

- Active Power (kW)
- Reactive Power (kVAr)
- Apparent Power (kVA)
- Voltage (V)
- Current (A)
- Frequency (Hz)
- Power Factor

At the USU EVR facility, power monitors are used to collect data from the power system. A Leviton 3500 series monitor enables accurate measurements and continuous collection. Due to the familiarity and reliability of the Leviton 3500 series monitors, these are the power monitors recommended for the Intermodal Hub site, if other monitoring means cannot be achieved. Identified locations for monitoring areas are: Bus depot chargers, overhead chargers, Track and Power Substation (TPSS), snowmelt system and CNG station. These locations are shown in Figure 11.

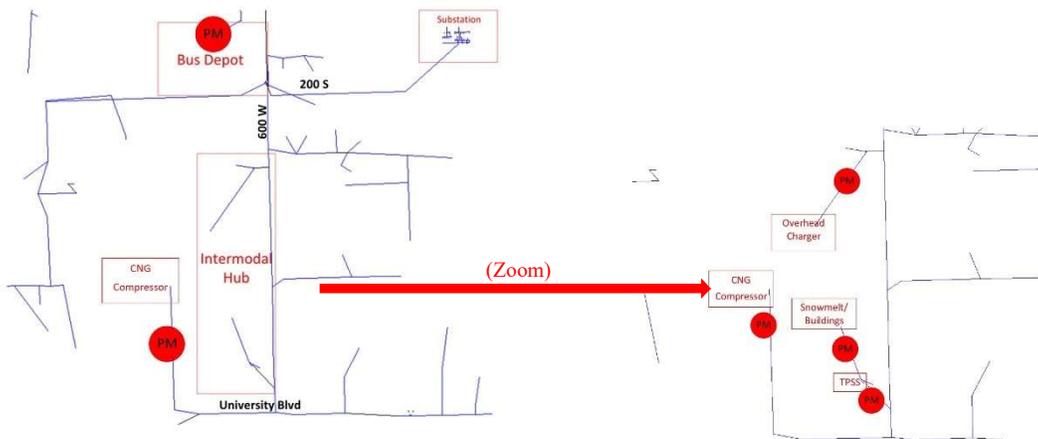


Figure 11 – Proposed Power Monitor Installations

There are multiple ways to collect the data from the power monitors. Many communications

protocols are used for data collection; CAN bus, Modbus, DNP, etc. The monitors used at the USU EVR facility utilize Modbus TCP/IP communication protocol. Modbus TCP/IP utilizes a master client format and is similar to Modbus RTU. The major difference between the two is that TCP/IP interface runs over Ethernet. Due to the familiarity of Modbus TCP/IP, this will be the communications protocol used to collect data from the monitors at the Intermodal Hub site. In order to utilize Modbus TCP/IP and to have instantaneous data, the monitors need to have a network connection. Figure 15 shows an example of how devices are connected using Modbus TCP/IP.

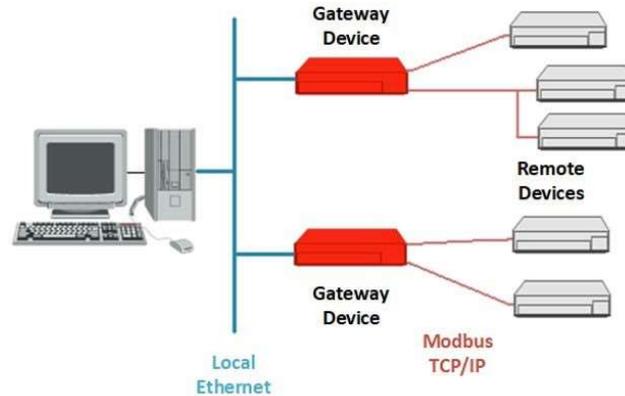


Figure 12 – Network Diagram

At the USU EVR facility, the power monitors have a hard-wired Ethernet connection which is connected to the facility network. Due to infrastructure limitations, it is not possible to have hard-wired connections at the Intermodal Hub site. Therefore, monitors installed at the UTA site will be connected to a Verizon cellular network. This requires additional hardware for the Intermodal Hub site. In order to transmit the data over a cellular network, a router must be installed between the meter and the device requesting the data. The router will need to accept a Verizon sim card and allow the devices IP address to be static. This will ensure that IP addresses assigned to the monitors, routers and sim cards will not change. After reviewing the hardware requirements for connecting to Verizon’s network, the Cradlepoint IBR650C meets those requirements. Figure 13 demonstrates how the data collection system at the Intermodal Hub site will operate.

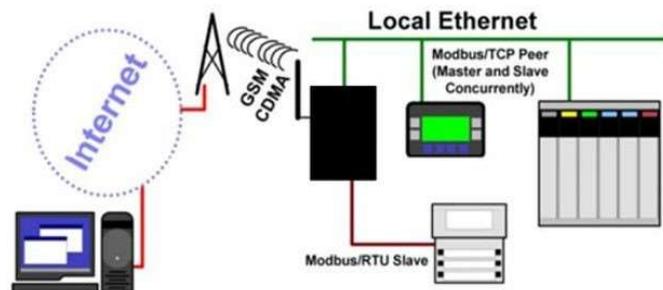


Figure 13 – Proposed Data Collection System

In order to maintain data privacy and prevent vulnerabilities, USU has worked with Verizon to establish a zero-tunnel private network for data transfer at the Intermodal Hub site. A zero-tunnel

private network segregates the devices at the Intermodal Hub site from the public internet while enabling devices to communicate over IP addresses. The network is only accessible by a private network gateway and devices with authorized IP addresses. This ensures data safety and reliability of devices connected.

Transferring the data from the power monitor to the algorithm is the next challenge. To do it properly and meaningfully a two-server approach will be used to solve the challenges. The primary benefit of using two servers, is that we can use one server as back up in case of server failure in our main server. Another reason for opting for the two servers is that the work has been divided into two separate machines and it decreases the chance of total failure.

The following Figure 14 shows conceptually how this communication system will function. The primary components of this system are: data collection server, data pre-processing server, algorithm, and power monitor meters.

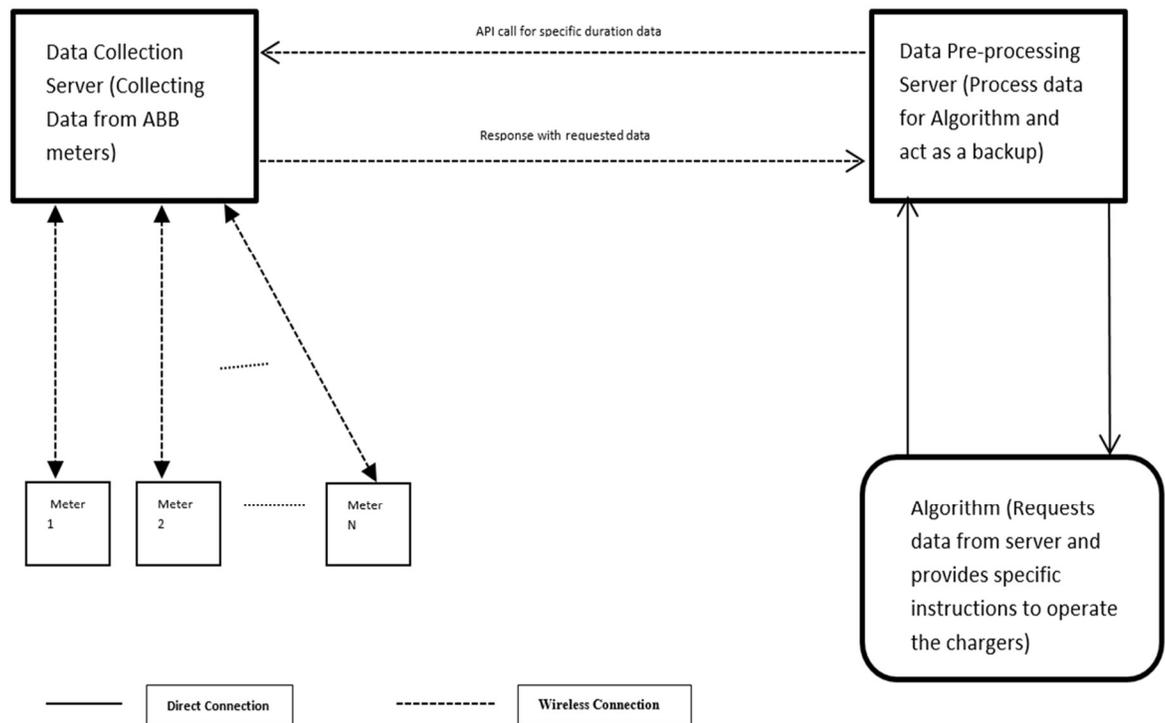


Figure 14 – Communication Architecture

A description of each of these components follows:

Data Collection Server stores response from meters and provides the data to the data-preprocessing server for the desired time period. It also transfers the instructions from the algorithm to specific chargers for action.

Data Pre-processing Server communicates with both the data collection server and our Reinforcement Learning algorithm. It collects data from the data collection server and pre-processes the data according to the algorithm requirements. It receives the algorithm output and processes them into machine instructions.

The algorithm will only communicate with the data pre-processing server over a dedicated

connection. The algorithm requests data for a specific period and processes this data to develop action sequences for the charging units.

Meter sends charge data only to the server at a defined frequency.

The meter to Server connection is a dedicated communication pathway for the meter to transmit readings to the server. The pathway is a zero-tunnel private network mentioned above.

For the data collection server to data pre-processing server connection, a wireless connection with Application Programming Interface (API) call will be implemented. For security, the data server will only send data to the data pre-processing server. For the wireless connection, we will use the internet to communicate between the servers. Both servers are proposed to be a web host and use an API call, specific to the servers, to send data or request data from one server to another.

Server to Algorithm connected with a dedicated wire connection to reduce the security issue and latency.

3.2 **Write code and program algorithms on servers to include energy/load balancing and management**

Algorithm Concepts: Planning and Operations at the EVR

The algorithm development described in Task 2 characterizes smart load balancing and smart charging as a single problem for reinforcement learning. For the EVR deployment, these problems are addressed using a two-level hierarchical strategy consisting of planning (level 1) and real-time operation (level 2). Reinforcement learning may be applied in the operational level. These levels are described below.

Level 1: Planning

Planning is a process that receives information about the system to be controlled and outputs a plan, which is a sequence of decisions for operating the system in an optimal manner at each instant in time. Planning is carried out over an appropriately chosen time horizon. In the setting of bus fleet operation at the UTA, the natural time horizon is one day. Each day the system is operated according to the same plan. A separate plan would be derived for weekend days when bus routes differ from weekdays.

Planning relies on determinism. Planning considers an ideal day and assumes that buses follow their routes and are on time. It assumes that electric bus batteries discharge along their routes in a manner that obeys dynamic models. Uncontrolled system loads (such as TRAX trains) are assumed to be known in advance and are perfectly predictable. In its most basic form, planning does not consider variations due to traffic delays or weather changes or variable ridership. Everything important about the system is assumed known in advance and the system is mathematically predictable.

Planning solves an optimization problem. Given knowledge about the system to be controlled, planning establishes a set of unknown variables related to the sequence of actions to take. Using problem data, an extensive set of constraints are derived. These constraints require, for example, that bus batteries retain sufficient charge to return to the station, enforce conservation of energy at electric meters, and enforce physical dynamics of bus motion. An objective function measures the peak load and energy consumption, which are to be minimized. The

resulting problem may have 100,000 or more variables and constraints. Using the tools of numerical optimization and modern computers, the problem is solved, and the solution variables represent the plan. The plan indicates when each bus should charge at which station and for how long. At every time instant throughout the day, the plan characterizes the battery states of charge for every bus in the fleet. The plan answers questions about what the minimum peak load will be and how to achieve it.

The problem with planning is that it is too rigid. The real world does not obey the assumptions underlying planning. There are traffic delays and detours, weather variations, and unaccounted for spikes in ridership and other electric loads at the Intermodal Hub site. Planning requires that everything be known in advance, but no actual day ever follows the assumptions of the ideal day used in planning. Planning can be extended to incorporate uncertainty by modifying constraints, but this only leads to a more conservative plan.

Level 2: Operation

Real-time operation addresses the practical problem of operating the system forward in time from the current state of the system. The plan might be that at some time instant Bus X should have 70% battery state of charge (SOC) but instead it has only 50%. In practice, every bus battery will have a different SOC and be in a different location than what was planned. External loads will also differ from what was supposed to have happened on the ideal day.

When the system is found to reside in a state that differs from the plan, what should be done? One approach would be to return to a planning process. This would involve setting up and solving a new optimization problem with the current state as the initial condition. With 100,000 variables, re-solving the optimization problem is time consuming. Even if sufficiently fast computers could be applied to the planning problem, solving for a new plan whenever the state of the system deviates from the previous plan is not needed because there are alternatives that offer good performance and have much lower complexity.

One approach is like planning in that it involves optimization, but it focuses on a short time horizon. Planning ahead for 15 or 30 minutes reduces the problem to a manageable size so that it can be re-solved every few seconds. The objective used for operation is to move from the current state back to the optimal plan 15 or 30 minutes in the future. Thus, the objective function in the optimization for real-time operation is to measure the deviation of the current state from the optimal state 15 or 30 minutes ahead. Recognizing that the actual state of the system deviates from the ideal planned state, real-time operation attempts to return the state of the system to the state for the optimal plan. By applying this operation continuously, any deviation or upset can be accommodated and the operation loop continuously monitors and restores the system to the optimal plan.

Another approach to real-time operation applies reinforcement learning instead of numerical optimization to return the state of the system to the optimal planned state. This approach involves extensive training using simulated data at first and then switching to real-world data. Over time the agent learns to recognize appropriate actions to take given the actual as well as the optimal states by minimizing the penalty (negative reward) which is the discrepancy between the current state and the desired state.

Algorithm Implementation: Modeling, Simulation, and Optimization at the EVR

This section describes the current status of the level 1 (planning) and level 2 (operation) algorithms at the EVR.

Planning

We have cast the planning problem as a network flow problem on a directed graph. The graph is an arrangement of nodes in a rectangular grid with one row for each bus and one column for each time step. The presence or absence of nodes in the grid indicates whether a bus is located at the station (node present) at the station at the given time. Otherwise, there is a “hole” in the grid of nodes at that time. Nodes in the graph represent the assignment of chargers to buses. A “non-charging” row of nodes indicates that a charger is not connected to any bus.

Nodes in the graph are connected by edges with a binary zero-one variable assigned to each edge. An edge variable with value one means that the edge is part of a path through the graph, whereas an edge variable with zero value means that the edge was not selected to be part of a path. An edge between two nodes on a bus row indicates that the charger was connected to the bus during those two time intervals. An edge between a non-charging node and a bus node or vice versa indicates a charger-to-bus connection or disconnection. The work done by a given charger to charge bus batteries is represented by a path through the graph. The edge variables with value one encodes the sequence of actions to charge bus batteries. The planning optimization solves for these binary valued variables.

Constraints are added to the planning optimization to reflect physical realities. The nodes representing all the times steps when a particular bus visits the station are members of a node group. Constraints on edge variables to prevent paths from visiting a group multiple times reflect the practical idea that the bus will connect to charging at most one time during a single visit to the charging station. Other constraints are derived from discretized dynamics of bus battery discharge and charge behavior.

The objective function is derived from a cost associated with each edge. If a given edge is selected (edge variable with value one), then the associated edge cost is accumulated into the objective function. In the current problem formulation, edge costs account for energy consumption.

With all of the preceding matters in place, the planning optimization problem takes the form of a mixed integer linear program (MILP). This is a common type of problem and many commercially available solvers are available for this problem. To date we have used the Gurobi optimization package and called it from Matlab. The problem size scales with the granularity of the time quantization. In one instance of this problem, we used a one-minute time interval, 24 buses, and three chargers. For a one-day time horizon, there were 146,000 edges and 16,000 charging variables. On a standard laptop, a feasible solution was found in 10 minutes and an optimal solution was found in 100 minutes. That seems like a long time, but it should be recognized that the plan need be computed only once. Clearly, recomputing the planning solution is not possible in a real-time operational scenario.

Our problem formulation improves upon other methods which assume that an electric bus charges for the entire interval while it visits the charging station. Our network flow formulation optimizes a plan to use charging only when it is needed during a visit to the station. The current formulation still needs to be extended to be applied at the EVR in Logan or the Intermodal Hub site in Salt Lake City, and we are currently working on the following extensions.

- Minimization of peak load. We are currently working to incorporate minimization of peak load into the network flow-based optimization.
- Continuous vs. bang-bang charging. The version of planning described above charges buses at either zero (no charging) or the maximum rate. This is known as “bang-bang”

charging because the charging rate toggles between off and maximum on rates. However, the real ABB chargers and the OCPP protocol allow charging on a sliding scale. We are currently working to incorporate continuously variable charging into the network flow optimization.

- **Battery health.** A third extension that is currently being studied is accounting for battery health. Cycling, elevated temperature, and aging decrease battery performance and capacity over time. By managing discharge-charge cycling frequency and depth, a Lithium-based battery can deliver up to three times the energy over its lifetime. Thus, accounting for battery health in making charging decisions can significantly prolong battery life for the bus fleet operator and save a lot of money.

Our early work on optimization for planning appears in the paper “A Network Flow Approach to Battery Electric Bus Scheduling” by Whitaker, Hansen, Droge, Mortensen, and Gunther. The paper is currently in draft form.

Real-Time Operation

Real-time operations are guided by an optimal plan. Our planning work is not mature enough to be coupled with operational optimization. Therefore, our work on real-time operation has proceeded using an instantaneous opportunistic plan. This approach substitutes an optimal plan with one derived from the current state of the system. This is a reasonable approach but guarantees only local optimality rather than global optimality. The EVR demonstration is based on the opportunistic plan approach.

Operations with opportunistic planning is divided into two stages: smart peak management and smart charging. The peak manager observes the power consumed by the non-controlled loads compared to the current running peak 15-minute average power usage. The difference between these two is the power budget available for charging bus batteries that avoids raising the peak. Smart charging decides how to allocate the available budget among the chargers with two objectives of (1) not raising the peak average power and (2) charging bus batteries to their optimal charge. Normally a plan would provide a target battery charge level. In opportunistic planning, the optimal battery charge level is determined by considering the distance the bus must travel before its next charging opportunity. The difference between the target charge level and the battery’s current state of charge divided by the time to depart is the optimal charging rate. This is computed for each bus and the rates are calibrated based on real bus data collected through the New Flyer web interface. The actual charge rate selected for each bus is in proportion to the sum of the charge rates for the buses connected to all the chargers multiplied by the power budget.

The baseline for comparison uses no optimizations. The baseline simply charges buses at a fixed charging rate while the bus is at the station or until full battery charge is reached.

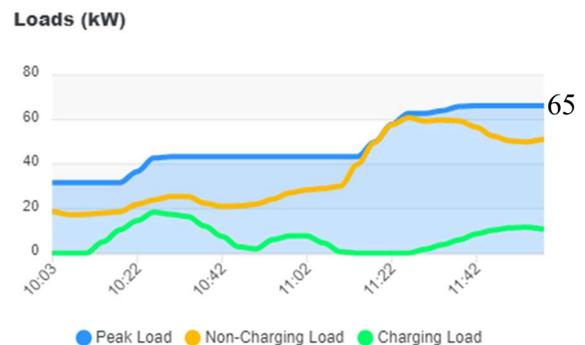
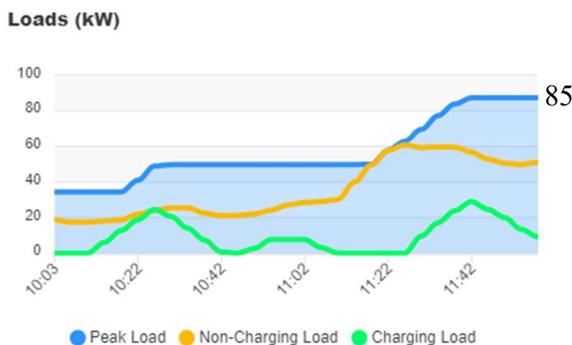


Figure 15 - Comparison of Standard (left) and Smart (right) Peak Management

A comparison of standard and smart peak load management appears in Figure 15. The green line is the power used to charge buses. The yellow line is the power used by non-controlled loads (everything else except for bus battery charging). It is the same in both graphs. The blue line is the 15-minute peak average power used. This line illustrates how smart peak management reduces peak demand. When enabled, smart peak management reduces the power budget available for charging when there is a rise in the power used by non-controlled loads (yellow line). The reduction in power used to charge buses is seen by comparing the green lines toward the right of the plots. The important feature is that, in this example, the peak average power (blue line) is about 20 kW lower for smart peak management than when no peak management is used.



Figure 16 - Illustrates Smart Peak Management Combined with Smart Charging

In Figure 16, the combined effects of both smart charging and smart peak management are illustrated. The dark and light blue lines represent the peak average power and battery charger power, respectively, when there are no optimizations enabled (baseline). The dark and light red lines show the same quantities when both peak and charging optimizations are enabled. As in the previous example, the savings in peak average power is approximately 20 kW. The light blue and light red lines illustrate that the optimized charging distributes the charging budget more uniformly over the seven-hour time period. This more intelligent reassignment of charging resources has a dramatic effect on peak demand while keeping the buses fully operational. Other examples have been saved in movie format for real-time viewing and are accessible via an archive separate from this report.

3.3 Evaluate hardware system (with integrated software) at the USU EVR

Controlling ABB Chargers via Server

ABB chargers can be managed via a server using Open Charge Point Protocol (OCPP), a technology that provides a universal way to regulate electric vehicle chargers. At USU we tested the capabilities of a server meant to control multiple ABB chargers. In doing so, we discovered that ABB chargers support a portion of the features OCPP is capable of, including charge power limitation, data collection, and charger configuration.

Capabilities Tested:

Boot Notification

Notification that displays when charger is powered on and connected. The information received through boot notification is charge point vendor, charge point model, charge point serial number, and firmware version. Among the information, we can use the charge point serial number to identify the charger uniquely. The figure below is a response from the boot notification function.

```
Boot Notification: MD_HVC_CAR ABB {'charge_point_serial_number': 'T175-IT1-3420-016', 'firmware_version': '1.3.5.4'}
```

Figure 17– Boot Notification Response

Heartbeat

The heartbeat function is called from the ABB chargers in the certain interval set by the server. It keeps track that the charge point relates to the OCPP server. If the heartbeat connection is lost, connection retry is initiated to the charge point from the server. The only value it provides us is the charging point time.

Status Notification

Notification that displays the status of the chargers (i.e. available, charging, preparing, etc.). In the status notification, we are getting the values of connector id, error status, the actual status of the connector id, time. Every time the status changes for the connector it reports to the server.

```
From the status notification: 0 NoError {'status': 'Available', 'timestamp': '2021-03-13T00:34:55.123Z'}
From the status notification: 1 NoError {'status': 'Charging', 'timestamp': '2021-03-13T01:11:04.549Z'}
From the status notification: 2 NoError {'status': 'Available', 'timestamp': '2021-03-13T00:35:13.331Z'}
```

Figure 18 – Status Notification Response

Change/Get Configuration

Changing and displaying the configuration/settings of the charger. From the get notification, we can see the ABB charger configuration and the corresponding values. The configurations include interval time of meter values, heartbeat, and simple clock data. The values we want to keep track of in an active charging session are also in the configuration. Here is the ABB configuration we are getting from the get configuration call below.

```

ConnectorPhaseRotation: Unknown
LocalAuthListMaxLength: 10000
ChargeProfileMaxStackLevel: 10
LocalPreAuthorize: false
GetConfigurationMaxKeys: 200
AllowOfflineTxForUnknownId: false
LocalAuthorizeOffline: false
WebSocketPingInterval: 30
SupportedFeatureProfiles: Core,LocalAuthListManagement
ChargeBoxName: DefaultChargerName
MeterValuesSampledData: SoC,Energy.Active.Import.Register
AuthorizeRemoteTxRequests: false
ConnectionTimeout: 30
NumberOfConnectors: [2]
ResetRetries: 5
ContactCenter: Default contact center details
AuthorizationCacheEnabled: false
ChargeBoxTimeZone: +0
StopTxnAlignedData:
SendLocalListMaxLength: 1000
StopTxnSampledData:
StopTransactionOnInvalidId: true
LocalAuthListEnabled: false
TransactionMessageAttempts: 10
MaxChargingProfilesInstalled: 10
ChargingScheduleMaxPeriods: 24
UnlockConnectorOnEVSideDisconnect: true
ClockAlignedDataInterval: 20
TransactionMessageRetryInterval: 60
ChargingScheduleAllowedChargingRateUnit: Current,Power
ReserveConnectorZeroSupported: false
HeartbeatInterval: 10
MeterValuesAlignedData: Temperature
MeterValueSampleInterval: 15
MinimumStatusDuration: 5
StopTransactionOnEVSideDisconnect: true
FallbackDelay: 300
FallbackMode: NO_FALLBACK
FreevendEnabled: true
FreevendIdTag: NOA

```

Figure 19 – Get Configuration Response

Using the change configuration function, we can change the ABB configuration from the server.

Reset

We can initiate a charger reset from the server using the reset function. The two options for resetting are hard and soft. A hard reset is a last resort function that forces processes to stop and restart. A soft reset is safer and is more commonly used to implement changes after modifying charger settings.

Meter Values

Meter values allow for the collection of data both while the vehicle is charging and while the charger is inactive. An interval can be set so that values are sent regularly. In figures 21 and 22, the time interval is set to 15 seconds which is why the charge limitation is displayed after 15 seconds. Using the change configuration feature, meter values can be modified to display multiple statistics including, but not limited to current, vehicle battery percentage, voltage,

temperature, and power.

Start Transaction

This is a notification that sends when a vehicle begins charging. It includes information about the vehicle and charger including energy status, vehicle ID, and time.

Set/Clear Charging Profiles

Capabilities include adding and removing charging profiles that store charging limits, charge start times, etc. This allows for the ability to control charge power/current at the start of and in the middle of a charging session.

Below is an example of a meter value response after limiting the power of a charger. The underlined values are the power limit in Watts and the time in UTC. The following figure is the charger statistics before the limitation is sent. As shown in the underlined values, the vehicle is charging at 13000Wh at 1:12:49pm.

```
Meter values called.....
Meter values: 1 [{ 'sampled_value': [{ 'unit': 'Percent', 'context': 'Sample.Periodic', 'measurand': 'SoC', 'location': 'EV', 'value': '75.0'}, { 'unit': 'Wh', 'context': 'Sample.Periodic', 'format': 'Raw', 'measurand': 'Energy.Active.Import.Register', 'location': 'Outlet', 'value': '481147'}, { 'unit': 'W', 'context': 'Sample.Periodic', 'measurand': 'Power.Active.Import', 'value': '15621'}, { 'unit': 'W', 'context': 'Sample.Periodic', 'measurand': 'Power.Offered', 'location': 'Outlet', 'value': '13000'}, { 'unit': 'A', 'context': 'Sample.Periodic', 'measurand': 'Current.Import', 'value': '41'}, { 'unit': 'A', 'context': 'Sample.Periodic', 'measurand': 'Current.Offered', 'location': 'Inlet', 'value': '20'}], 'timestamp': '2021-03-13T01:12:49.627Z'}] { 'transaction_id': 1}
```

Figure 20 – Meter Value Before Charge Limitation

The next figure shows the meter values after the charging profile is changed while the vehicle is charging. The underlined values show the charge cap is changed to 9800Wh at 1:13:04pm.

```
Meter values called.....
Meter values: 1 [{ 'sampled_value': [{ 'unit': 'Percent', 'context': 'Sample.Periodic', 'measurand': 'SoC', 'location': 'EV', 'value': '75.0'}, { 'unit': 'Wh', 'context': 'Sample.Periodic', 'format': 'Raw', 'measurand': 'Energy.Active.Import.Register', 'location': 'Outlet', 'value': '481186'}, { 'unit': 'W', 'context': 'Sample.Periodic', 'measurand': 'Power.Active.Import', 'value': '11430'}, { 'unit': 'W', 'context': 'Sample.Periodic', 'measurand': 'Power.Offered', 'location': 'Outlet', 'value': '9800'}, { 'unit': 'A', 'context': 'Sample.Periodic', 'measurand': 'Current.Import', 'value': '30'}, { 'unit': 'A', 'context': 'Sample.Periodic', 'measurand': 'Current.Offered', 'location': 'Inlet', 'value': '15'}], 'timestamp': '2021-03-13T01:13:04.627Z'}] { 'transaction_id': 1}
```

Figure 21 – Meter Values After Charge Limitation

Limitations:

- All the OCPP functions are not yet implemented by ABB, including the ability to create a schedule and reserve chargers.
- OCPP documentation of function names, variable names, spacing rules, and data types are frequently different than what is used with the ABB chargers.
- Only one version of OCPP (OCPP 1.6) is implemented. This is a drawback because OCPP 2.0 includes many helpful features including the ability to send custom messages to the charger display.

- Some part of adjustment needs to call another function to activate. For example, many functions require the charger to be reset, but there is no documentation explaining when this is required.
- Some specific issues are still being sorted out by ABB. An example of this is when ABB updated their data communication from 3G to 4G, which broke the connection between our server and the charger until we worked through the issue with ABB.
- Errors that occur in the charger cannot be seen from the server. Only ABB staff can view these, which makes debugging and testing difficult.

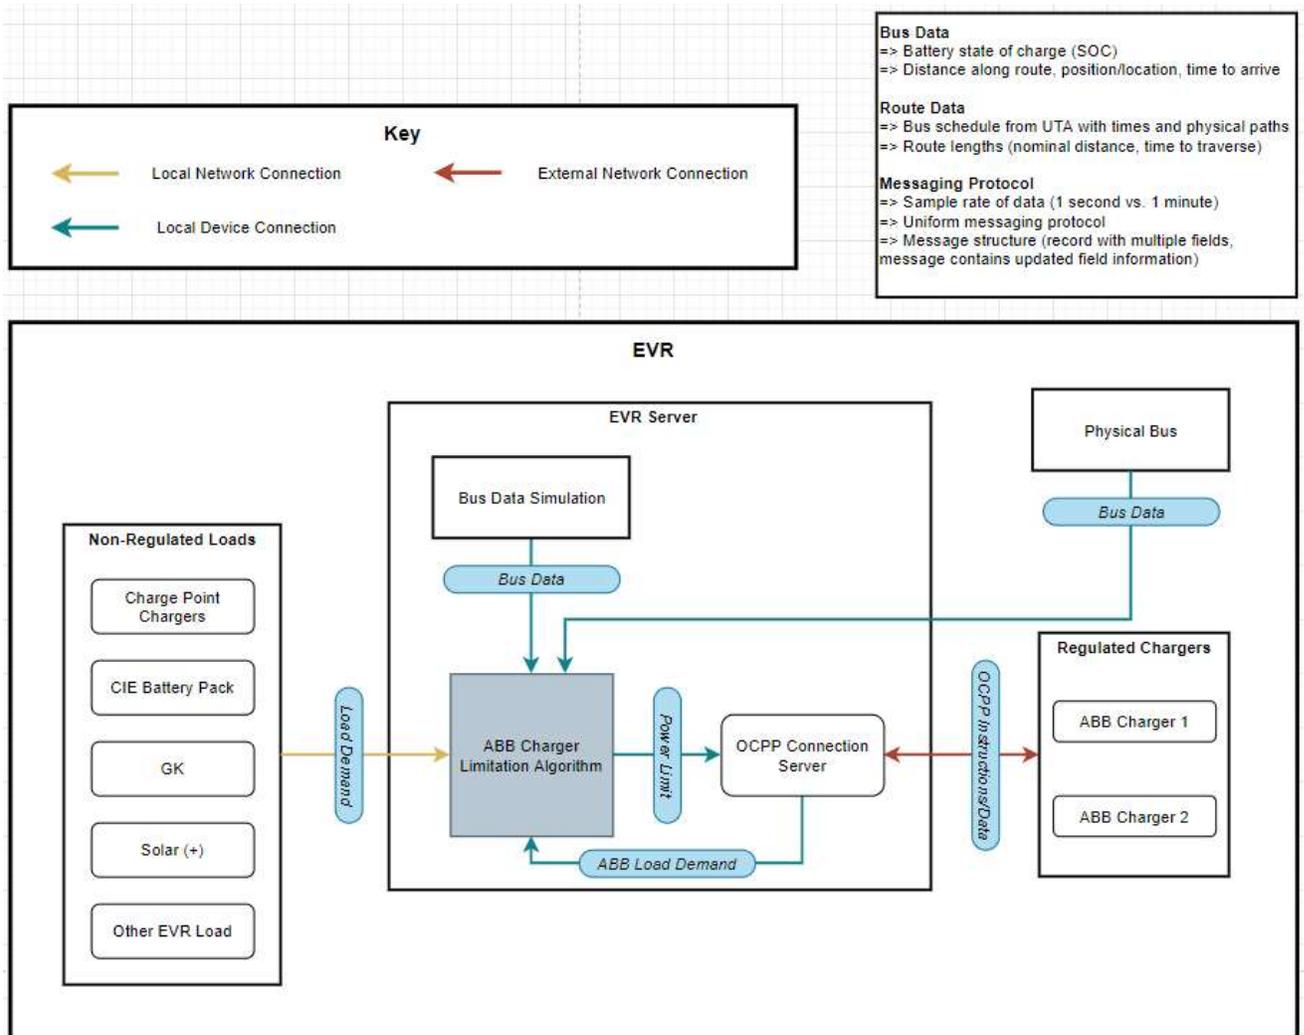


Figure 22 – EVR Testbed Diagram

Bus Data Simulation:

This part of the test consists of bus route information, SOC changes, and route time estimation.

Bus SOC change information:

We have used the new flyer API to simulate the state of charge (SOC) changes for four UTA

routes: 2, 6, 205 and 509. At the time of our initial testing electric buses only operated on route 2, and the New Flyer API was utilized to calculate the SOC change based on distance and elevation for that route. This data was used to create a model that can be applied to the other 3 routes that did not operate electric buses (6, 205, 509). Data from the UTA API was used from the other 3 routes, such as distance and elevation, to create a simulation that estimates SOC change, using the model from route 2. Figure 23, shows a portion of data received from the New Flyer system.

```

From: TO U HOSPITAL
Start time: (17, 58)
Start SOC: 72.4
To : TO SL CENTRAL
Reached at Destination: (18, 19)
End SOC: 72.0

From: TO SL CENTRAL
Start time: (18, 43)
Start SOC: 82.4
To : TO U HOSPITAL
Reached at Destination: (19, 5)
End SOC: 78.4

From: TO U HOSPITAL
Start time: (18, 58)
Start SOC: 80.8
To : TO SL CENTRAL
Reached at Destination: (19, 24)
End SOC: 78.8

From: TO SL CENTRAL
Start time: (19, 43)
Start SOC: 78.4
To : TO U HOSPITAL
Reached at Destination: (20, 3)
End SOC: 74.8

```

Figure 23 – A snapshot of new-flyer data used in the algorithm

Route information:

The bus route information, for routes 2, 6, 205 and 509, were collected from the rideUTA API. Data collection frequencies were established to obtain GPS location, speed, direction, line reference and the route and distance in a round trip, from the vehicles on this route. Since we have access to all this route data, we have been able to estimate average route timings, bus locations, time taken between stops etc. Figure 24 shows a sample of the collected data.

2021-05-24T14:53:36.26	2 TO U HOSPITAL	0.11508	40.76577	-111.91	84.4
2021-05-24T14:53:46.213	2 TO U HOSPITAL	0.5754	40.76577	-111.91	84.4
2021-05-24T14:53:56.61	2 TO U HOSPITAL	2.3016	40.76577	-111.91	84.4
2021-05-24T14:54:06.25	2 TO U HOSPITAL	6.44448	40.76582	-111.909	84.4
2021-05-24T14:54:16.3	2 TO U HOSPITAL	5.63892	40.76584	-111.909	84.4
2021-05-24T14:54:36.923	2 TO U HOSPITAL	2.64684	40.76581	-111.909	84.4
2021-05-24T14:54:47.303	2 TO U HOSPITAL	7.94052	40.76583	-111.909	84.4
2021-05-24T14:54:56.447	2 TO U HOSPITAL	4.25796	40.7652	-111.908	84.4
2021-05-24T14:55:06.46	2 TO U HOSPITAL	6.3294	40.76516	-111.908	84.4
2021-05-24T14:55:16.437	2 TO U HOSPITAL	21.51996	40.76444	-111.908	84.4
2021-05-24T14:55:36.48	2 TO U HOSPITAL	9.55164	40.76406	-111.909	84.4
2021-05-24T14:55:47.057	2 TO U HOSPITAL	7.82544	40.76428	-111.909	84.4
2021-05-24T14:55:56.513	2 TO U HOSPITAL	8.51592	40.76454	-111.91	84.4
2021-05-24T14:56:07.043	2 TO U HOSPITAL	5.06352	40.76463	-111.91	84.4
2021-05-24T14:56:16.643	2 TO U HOSPITAL	7.71036	40.76439	-111.91	84.4
2021-05-24T14:56:26.6	2 TO U HOSPITAL	6.21432	40.76417	-111.91	84.4
2021-05-24T14:56:46.87	2 TO U HOSPITAL	3.79764	40.76367	-111.91	84.4
2021-05-24T14:56:46.87	2 TO U HOSPITAL	3.79764	40.76367	-111.91	84.4
2021-05-24T14:57:06.97	2 TO U HOSPITAL	0	40.76362	-111.91	84.4
2021-05-24T14:57:16.987	2 TO U HOSPITAL	0	40.76362	-111.91	84.4
2021-05-24T14:57:26.797	2 TO U HOSPITAL	0.11508	40.76362	-111.91	84.4
2021-05-24T14:57:36.993	2 TO U HOSPITAL	3.56748	40.76361	-111.909	84.4
2021-05-24T14:57:56.797	2 TO U HOSPITAL	5.86908	40.76312	-111.909	84.4
2021-05-24T14:58:06.9	2 TO U HOSPITAL	7.36512	40.76313	-111.909	84.4
2021-05-24T14:58:17.037	2 TO U HOSPITAL	10.3572	40.76357	-111.909	84.4
2021-05-24T14:58:26.913	2 TO U HOSPITAL	9.43656	40.76397	-111.909	84
2021-05-24T14:58:36.85	2 TO U HOSPITAL	8.0556	40.76426	-111.909	84
2021-05-24T14:58:46.927	2 TO U HOSPITAL	7.82544	40.76445	-111.91	84
2021-05-24T14:59:07.293	2 TO U HOSPITAL	5.63892	40.76485	-111.91	84
2021-05-24T14:59:07.293	2 TO U HOSPITAL	5.63892	40.76485	-111.91	84

Figure 24 – Small part of route 2 data collection

Non-Regulated Load:

UTA has electric loads such as TRAX, snowmelt, and building infrastructure which are uncontrollable, meaning that the power schedule for these inputs cannot be controlled and will affect charging and discharging schedule for the buses. Therefore, our algorithm must make scheduling decisions of the influenced electric vehicles around these uncontrollable loads. For the EVR demonstration we did not have the capability to collect real-time data for these appliances yet, and instead used our EVR lab’s load to simulate these uncontrollable loads. Even though our lab’s load isn’t as high as the UTA loads, it will still exercise the algorithm to make some consideration for uncontrollable loads. Calculation of uncontrolled loads at the USU EVR include solar panels, AC, heating, generator, CNC machine, charge point chargers and other regular electric consumption. The graph below shows the load for around 5 days:

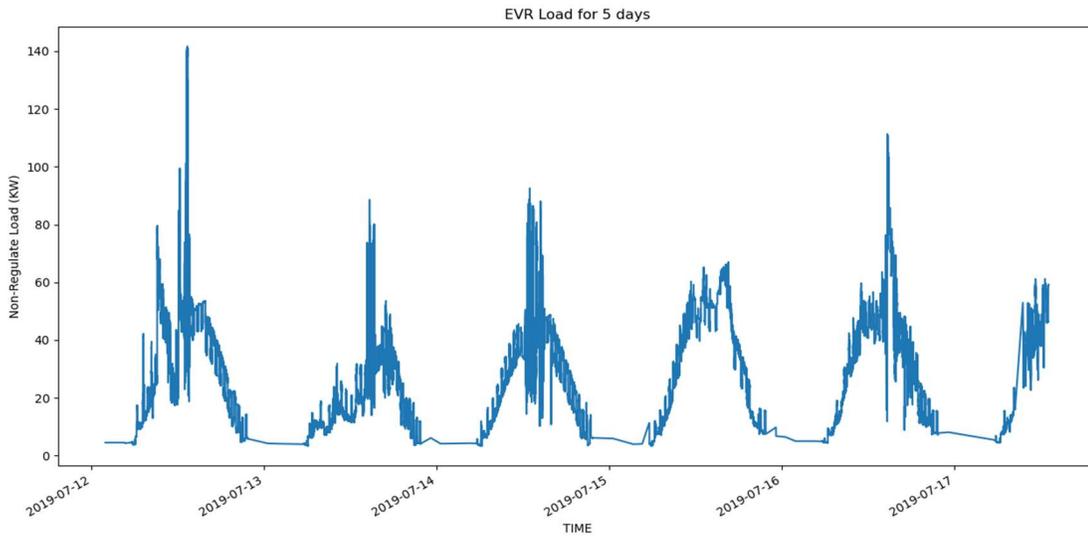


Figure 25 – EVR Load for 5 days

OCPP Instruction Verifications:

Through OCPP there is capability to control the ABB charge rate from the server in real-time. Output power can be dynamically adjusted or set from 0 kW to the maximum accepted by the vehicle. During the EVR demonstration, and with the absence of a bus, we were not able to demonstrate this charge setpoint control, however with the energy management system (EMS) we were able to demonstrate this. Examples of EMS results are depicted below, and provide control verification of the USU ABB high power chargers.

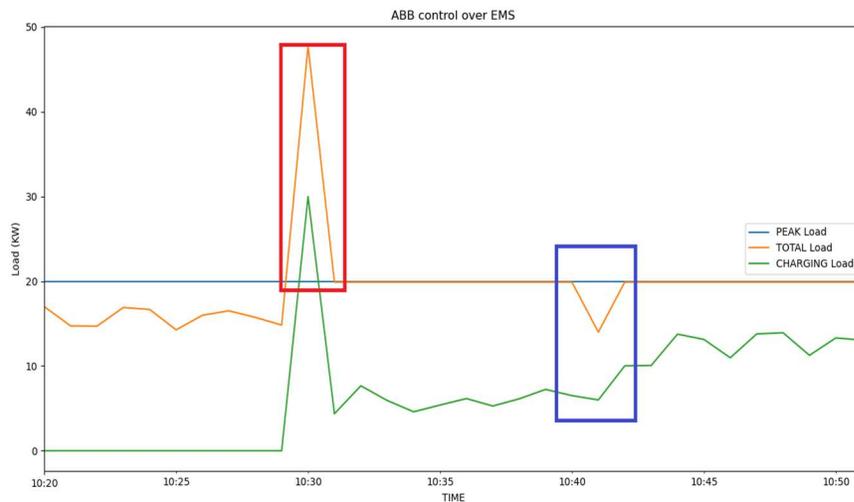


Figure 26 – EMS result for 02/22/2021

Our EMS is reflexively taking action every minute. The test target was to keep the EVR peak

below 20 kW. As can be seen by the green line, no charging load was experienced until approximately 10:20. No action was necessary by the EMS to maintain the target. However, at 10:30 a charge event occurred, eclipsing the target 20 kW. The EMS quickly responded lowering the charging speed to have the peak constant. The opposite effect can be observed, as highlighted by the blue box, when the non-controllable load went down and the EMS responded by allowing an increased available charge output.

UTA API:

Three UTA APIs were used to plan and develop the bus simulation. These APIs follow a standard data format for transit tracking called Service Interface for Real Time Information (SIRI). Data gathering from the APIs includes coordinates, speed, direction, and destination.

1. Vehicle Monitoring by Route

- Displays data for all buses along certain routes
- Used for calculating average wait times, finding the IDs of buses, and determining charging schedules

2. Vehicle Monitoring by Vehicle

- Displays data for a single bus
- Used for calculating route bus travel distance and route coordinates

3. Stop Monitoring

- Displays vehicles approaching a stop over a chosen time interval
- Used for calculating bus stop idle times

A combination of these APIs was used to develop maps for bus tracking, route length calculation, and bus availability for charging (Figure 27). These statistics were then used to develop a realistic algorithm that matches the real behavior of UTA buses based out of Salt Lake Central Station.

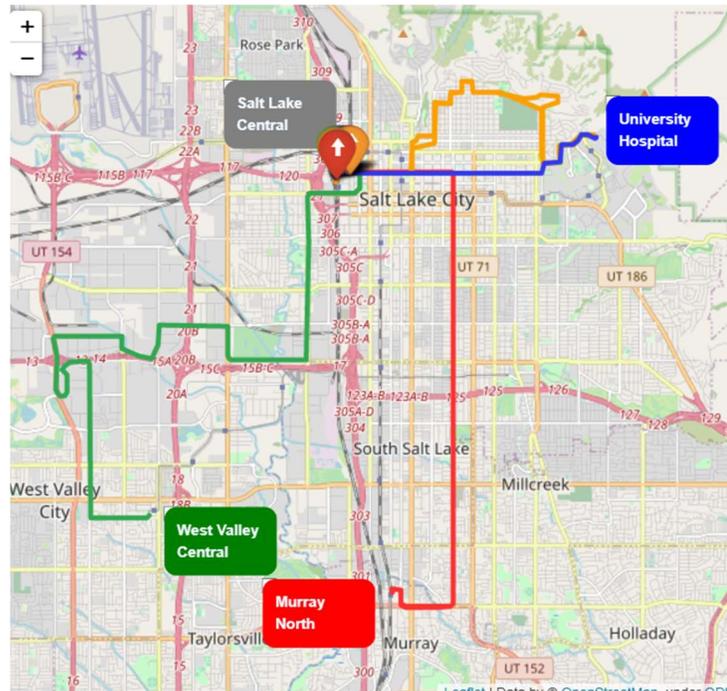


Figure 27 - Visualization of routes tracked using UTA APIs

3.4 Anticipate needs for and develop cyber security management

As wider demonstration of the algorithm is realized, discussion with UTA will be important to understand the cybersecurity considerations.

Since our algorithm functions on real-time route data, fake information can be used to disrupt our algorithm. We have found similar potential threats for the future:

- TRAX and snowmelt data - Faking high demands from TRAX may cause the algorithm to reduce charging speed for buses. Source verification should be considered for the TRAX information or use a secure VPN.
- Bus data – For this demonstration bus data was collected through liveUTA API. This may allow for fake bus information to be provided and result in a less optimal decision being made by the algorithm. Alternate data acquisition solutions could be implemented, such as data loggers and secure data collection paths.
- ABB chargers - ABB chargers use web sockets to communicate with servers. Using fake ABB chargers or tampering with existing ABB chargers might communicate wrong information to our algorithm. Additional solutions should be considered to verify the ABB charger information, such as someone physically confirming the information.

4 Task 4: Piloting and Field Evaluation

The goal of the Piloting and Field Evaluation task was to demonstrate a power balance and demand response system for vehicle charging at the Intermodal Hub site in SLC, which has high peak power demand. Our piloting and field evaluation shows that, compared to the default load management, smart load management reduces peak demand and increases utilization of the distribution system resulting

in prolonged use of existing infrastructure and reduced operating costs.

The piloting and field evaluation efforts were divided into multiple subtasks including: (1) accessing and exercising application programming interfaces, (2) software architecture for simulation and real-time control, (3) sliding-window planner for real-time control of vehicle charging, and (4) field evaluation. Each of these are discussed in the subsections below.

4.1 Accessing and exercising application programming interfaces

Real-time smart load management relies on sensing and control. Sensing is the acquisition of information about the current and historical state of the system to be controlled. Control signals are delivered to actuators to implement system commands. In the demonstration system, sensing and control are made possible through the application programming interfaces (APIs) described in this section.

The New Flyer API provides real-time information about the location of electric buses as well as the battery state of charge (SOC), battery temperature, and outside temperature. Watching battery SOC over time enables the construction of models of battery discharge for each route and enables the estimation of charge rate when buses are connected to chargers at the Intermodal Hub station as well as in the depot overnight. Both SOC discharge and charge rate vary with temperature. We are collecting data to build temperature dependence into the battery model in the future. The current deployment assumes batteries at a fixed temperature. Deviations due to temperature variations are accounted for through feedback control.

The UTA API provides information about bus location for all types of buses (electric, diesel, CNG). A database of bus travel paths over time provides the information needed for realistic simulation of a large number of electric buses before actually deploying those buses. Thus, data acquired through the UTA API lets us predict and visualize the performance of smart load management in the future as the bus fleet transitions toward a larger fraction of electric buses.

The location data (latitude and longitude) returned by the New Flyer and UTA APIs is translated into a binary representation indicating which buses are present at the station or away on their routes at any given time. Buses at the station are available for charging. Thus, by reformatting the location data in this manner, we are able to construct a mathematical graph that encodes charging actions that can increase load on the power grid. Our real-time planning algorithm optimizes the cost of power with respect to the variables in the graph to determine which buses charge at each time step.

The electric train (TRAX) consumes considerable amounts of power when accelerating to leave the station and consumes much less at other times. The “spiky” nature of this load has a strong impact on peak power demand. It is expensive to operate and can be the determining factor in the infrastructure power capacity. Charging buses without regard to the TRAX load increases power costs and stresses the grid infrastructure. Smart load management must account for TRAX and other uncontrolled loads. An interface to the TRAX Power Substation (TPSS) has been implemented. Through this interface, information about the TRAX load is tracked and responded to in real time.

The final API provides communication with the ABB charging equipment through the OCPP protocol. This API may be queried to determine the current battery SOC. This API is used to send control signals to charge the bus battery. This is the API where the control actions of smart load management are executed. This API is provided for each charger at the Intermodal

Hub station.

4.2 Software Architecture for Simulation and Real-Time Control

The internal planning architecture for smart load management has been implemented using the framework provided by the Robot Operating System (ROS) to enable near seamless transition between development using simulated or recorded data and deployment using real-time data. ROS provides a modular architecture in which safe parallelization of system processes can be achieved using an established publication-subscription (pub-sub) framework for inter-process communication. ROS also provides an ecosystem of tools for visualization and status monitoring, reconfiguring the planner or other components on the fly, and tools to manage build processes and code repository management for large projects with interdependent components. Overall, it facilitates short, parallel development cycles with testing and evaluation done seamlessly between simulated and real-time data. Furthermore, ROS runs on PC workstations for development purposes and can be installed and run in the production/demonstration hardware/server for the field evaluation.

The major benefit of the pub-sub framework provided by ROS is an inherent modularity of the architecture. It enables different components of the architecture to be swapped in and out with no modification to the remaining components. Each major component is known as a node of the system. Nodes accept defined inputs and produce defined outputs. The simulation architecture is shown in Figure 28. The planner node ingests the status of each bus as well as the latest SOC reported by the OCPP node. It can make requests for further information, such as bus route and TPSS information. The difference between the simulation and the real-time execution is “simply” the swapping out of the OCPP-Sim and API-Sim nodes for nodes that communicate with the actual OCPP interface and APIs. The data provided to the planner is equivalent and the planner node is unaware that it is communicating with a simulation or live system. This allows for the following configurations/modalities.

- **Virtual:** The virtual configuration is pure simulation and is depicted in Figure 28. Each bus is simulated as are the interfaces to each bus.
- **Live:** The live configuration uses real-time data with interfaces to the OCPP and bus APIs. Commands are issued to the OCPP for charging of particular buses and the bus status information is collected through a web interface to the UTA and New Flyer APIs. In reference to Figure 28, the OCPP-SIM and API Sim are replaced with actual interface nodes and the bus sim nodes are eliminated.
- **Constructive:** The constructive modality can actually consist of a number of different configurations. The most basic being the bus status data being played back from recorded data and the OCPP interface being simulated. However, the modularity of the architecture could easily facilitate a combination of real, recorded, and simulated busses with real and simulated interfaces for the OCPP.

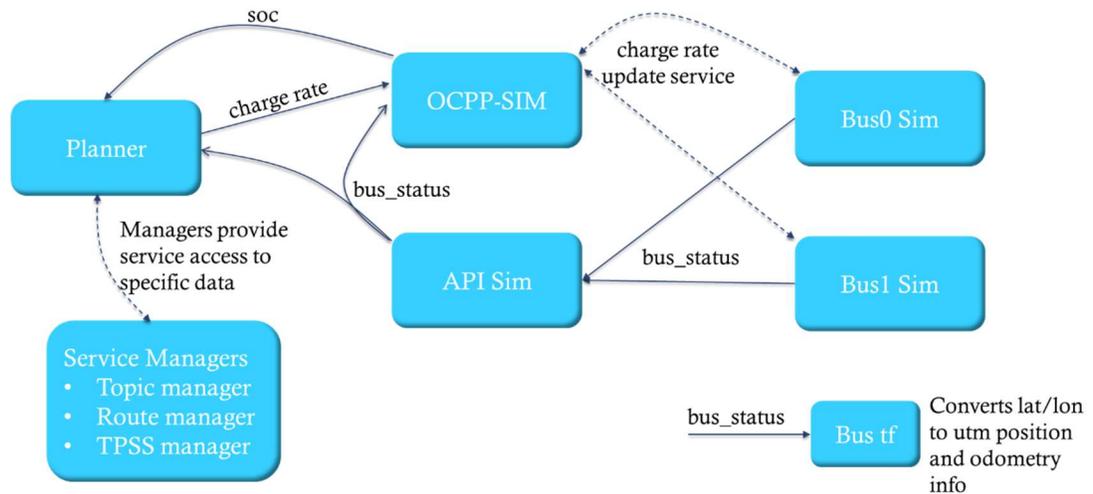


Figure 28 - Architecture for the Bus simulation. Nodes are depicted as rectangles, pub-interfaces are depicted as solid lines, and service calls are depicted as dashed lines.

Each of the components shown in Figure 28 has a specific role as described below.

- *Planner*: Ingests the status of each bus and creates a plan for when each bus should charge given the route information for the bus and its scheduled arrival at the bus depot.
- *Service Managers*: Provide historic and planned data for use by the planner and simulated components of the system. This data includes access to the route information for each bus, the historic data for the TPSS, and topics for pub-sub communication channels.
- *Bus tf*: Converts the spherical latitude, longitude, and altitude coordinates of the bus into a local cartesian frame for easier visualization of the location for each bus.
- *OCPP-Sim*: Simulates the interface to the OCPP. It monitors the status for each bus and accepts charge rates for each bus. For the live configuration, the OCPP-Sim is replaced by a ROS interface to the OCPP.
- *API-Sim*: Simulates data received by interfaces to the UTA and Newfler APIs. For the live configuration, the API-Sim is replaced with a node that periodically queries the UTA and Newfler APIs and encapsulates the data into a single ROS bus status message.
- *Bus-Sim*: Simulates the bus location and charge. These nodes are removed in the live configuration.

Each message and topic depicted in Figure 28 is defined in a ROS standard format and ROS tools are used to automatically generate C++ implementations to be used directly in the code. Most important is the *bus_status* message which contains a time stamp, an identifier for the bus, the state of charge, the temperature of the battery, the position of the bus, and the current activity, i.e., en route, at a station, or at the depot.

4.3 Sliding-Window Planner for Real-Time Control of Vehicle Charging

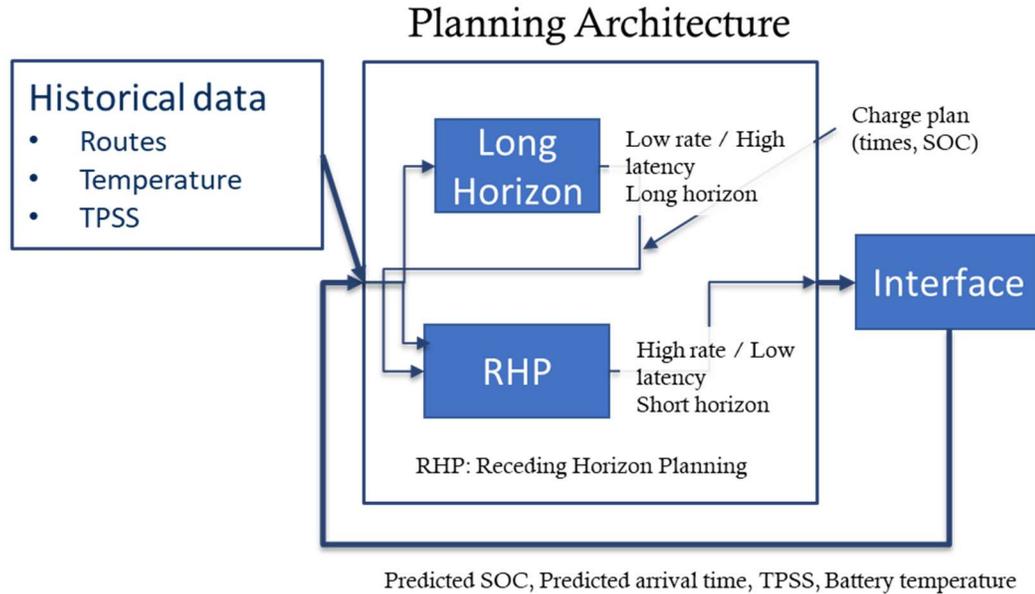


Figure 29 - Planning framework consisting of a long horizon and short horizon (receding horizon planner, RHP)

The AI plans for scheduling of bus charging using a hierarchical approach as depicted in Figure 29. The top-level planner develops an optimal charging schedule for a full 24-hour period under nominal conditions. These conditions utilize historical data for bus routes, temperatures, and TPSS power usage. The lower level planner uses receding horizon planning (RHP). This is a sliding window approach where the latest available data and predictions are used to adjust the plan over a certain period of time into the future.

A tradeoff exists on the length of the planning horizon on the look-ahead planner. While a very long look-ahead horizon for the sliding window planner would produce better plans, it also takes longer to compute and is, thus, not as adaptive to the latest information available. To allow for shorter horizons, an objective is included in the optimization to maintain consistency with the long-horizon planner. Note that for two buses with two overhead chargers, the long horizon planner can plan in less than a minute. Thus, for simple scenarios, the RHP is not needed. This is not the case when the number of buses and chargers is increased.

The optimization framework attempts to consider the total cost for bus charging. The cost is assumed to be a combination of three individual costs.

- *Energy or consumption cost*: The charge per kWh of energy consumed.
- *Facilities cost*: The maximum average power used in any 15 minute period of time.
- *Peak-time power cost*: additional cost on the maximum average power used during peak times.

The **energy cost** is a charge per kilo-Watthour (kWh) of energy consumed. There are different rates for on-peak times and off-peak times. Gain variables are used to encapsulate the energy in kWhs used during each discrete time period, this cost is easily captured in a linear cost vector associated with a mixed-integer linear programming (MILP) formulation. More specifically, the MILP formulation uses $\mathbf{y} = [\mathbf{x}^T \quad \mathbf{s}^T \quad \mathbf{g}^T]^T$ where \mathbf{x} contains the edge flow variables, \mathbf{s} contains the bus state of charge variables, and \mathbf{g} contains the gain variables. The

notation $g_{j,k,l}$ is used to denote the gain variable for bus j at time k with charger type l . The energy cost can then be written as

$$c_{energy} = \sum_k c_{energy,k} \sum_j \sum_l g_{j,k,l}$$

or, in vector form $c_{energy} = \mathbf{c}_{energy}^T \mathbf{y}$ where

$$\mathbf{c}_{energy} = [0^T \quad 0^T \quad c_{energy,0} \quad \cdots \quad c_{energy,N}].$$

The facilities cost and the power cost require additional considerations to formulate within a MILP framework for rapid calculation of optimal solutions. The **facilities cost** is defined as a cost on the maximum average power, in kilo-Watts (kW), used in any 15 minute period of time throughout the month. i.e., $c_{facilities} = c_f \max_t p_{15}(t)$ where c_f is the facilities cost multiplier and $p_{15}(t)$ is the average power over the 15 minutes prior to the time t . This can be written in matrix form by first using a selection and summation matrix, S , that sums the $g_{-}(j, k, l)$ for each time period. This matrix S is given by

$$S = [s_{q,r}] = \begin{cases} 1 & \text{if } q = k \\ 0 & \text{otherwise} \end{cases}$$

Then a Toeplitz structured matrix, T , is used to sum the energy for each time period and average it. This is done with

$$T = [t_{q,r}] = \begin{cases} \frac{1}{\Delta t} & \text{if } r \in \{q - m + 1, \dots, q\} \\ \frac{\delta \bmod \Delta t}{\delta} & \text{if } r = q - m \\ 0 & \text{otherwise} \end{cases}$$

By multiplying T and S by \mathbf{y} the average power is calculated for each time step. A slack variable is introduced to track the maximum of these average powers. This is accomplished with the constraint

$$T S \mathbf{y} \leq p_{max}$$

where the inequality holds for each row of T . This allows the facilities cost to be calculated in the MILP as

$$c_{facilities} = c_f p_{max}$$

The **power cost** is an additional cost on the maximum average power used during peak times. Thus, it is very similar to the facilities charge. The only difference is that the power cost only considers the average power over the on-peak times (e.g., 6-9 am and 6-10 pm). Thus, the power cost is calculated much the same as above with the addition of a time selection matrix, P ,

$$c_{power} = c_p p_{max,peak}$$

$$p_{max,peak} \geq P T S \mathbf{y}$$

$$P = [p_{q,r}] = \begin{cases} 1 & \text{if } q, r \in \text{Peak Times} \\ 0 & \text{otherwise} \end{cases}$$

The TRAX/TPSS load contributes to all three components of the rate schedule, but its main effect is felt in the peak demand. Figure 30 shows a plot of TPSS load over a 24-hour period. At this scale, the “spiky” nature of this load may be appreciated. However, zooming in to narrower time slices as in Figure 31(1 hour) and Figure 32 (6 minutes) enables visualization

of the short-term characteristics of TPSS. These figures illustrate that the TPSS load alternates between periods of consuming essentially zero power to periods where consumption is over 100 kW. The period widths appear random and the transitions between high and low load are nearly instantaneous. TPSS is an uncontrolled load. Because charging vehicle batteries adds on top of TPSS, the presence of TPSS creates a significant challenge for controlling vehicle charging while minimizing peak load.

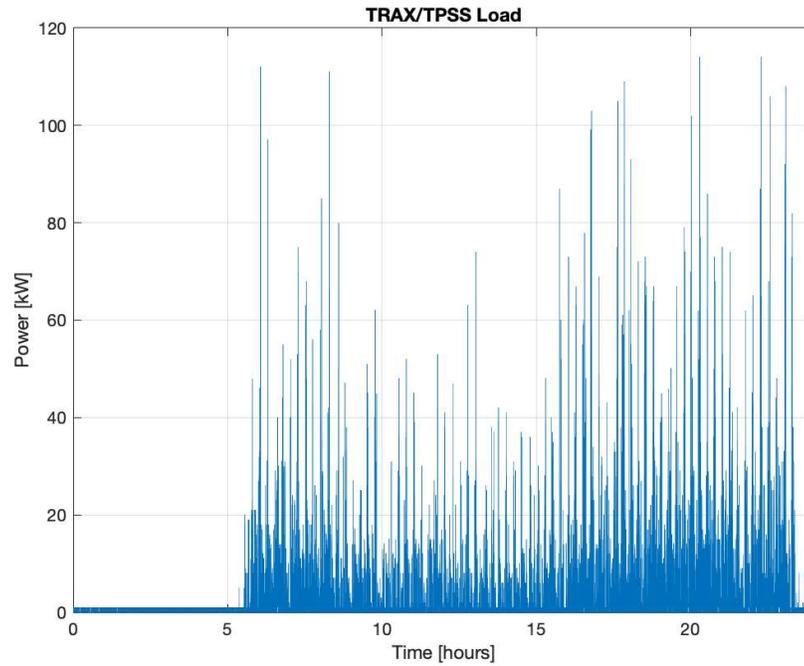


Figure 30 - Plot of TPSS load over one day (24 hours). The “spiky” nature of this load is difficult to visualize at this scale.

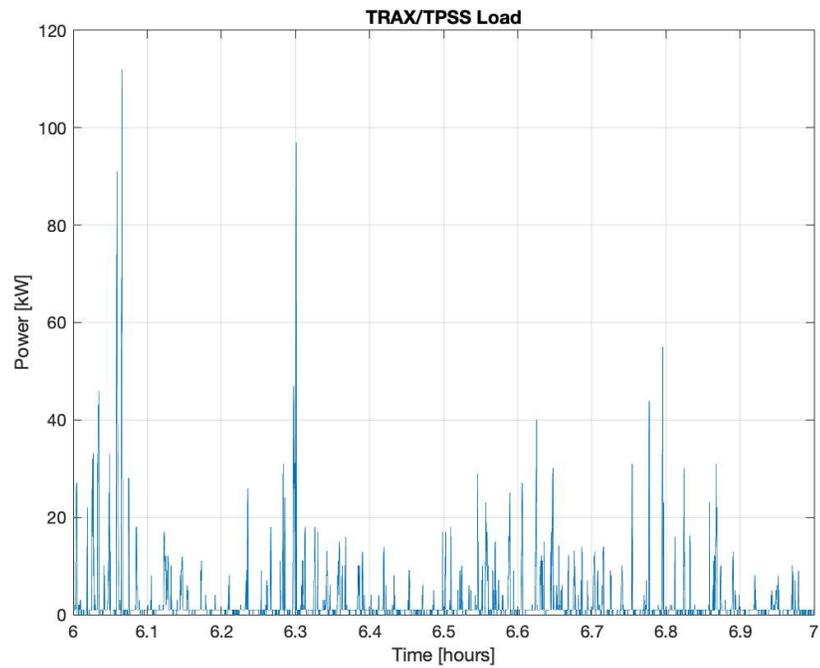


Figure 31 - TPSS load over one hour (6 – 7 AM). The “spiky” nature of this load is apparent

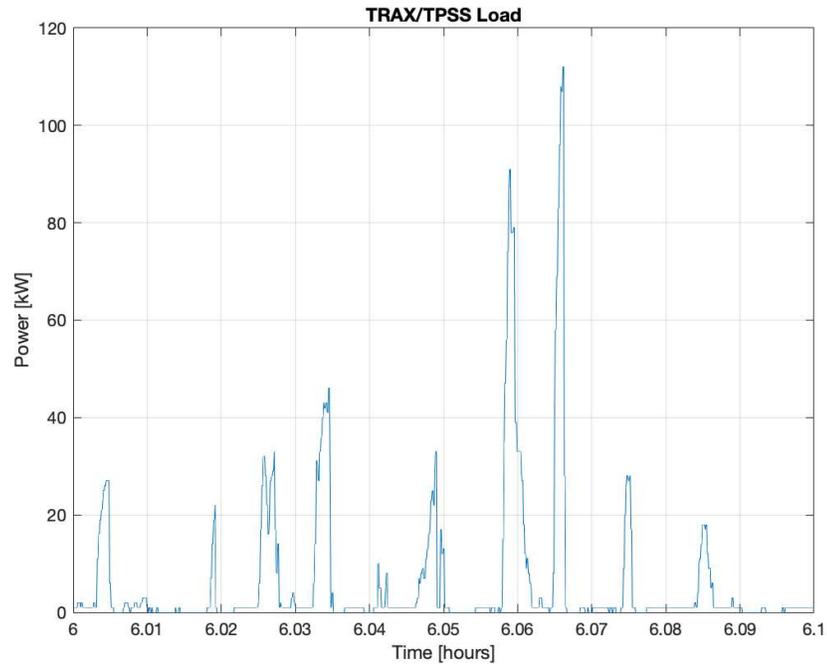


Figure 32 - TPSS load over 6 minutes. On this time scale, the intermittent structure of the TPSS load may be appreciated.

Modelling the SOC of a battery can be complex as it depends upon both the chemistry of the battery and the charging algorithm to be employed. The battery model alone can vary based upon temperature, current, and cycling. Care must be taken to balance the model accuracy with the ability to use the model within the optimization framework. A common battery charging control algorithm is CCCV: the charger applies a constant current (CC) to the battery until a battery terminal voltage is reached. The charger then holds that constant voltage (CV) as the charging current decreases. The charging current, which is approximately proportional to charge rate, is constant at low SOC and decreases as the SOC approaches full charge. We model this relationship using a first order differential system which, when given a step input, has a very similar exponential relationship between time and convergence to the maximum charge.

Historical data is used to model the expected drop in SOC of a bus along a particular route. This method allows the planner to make accurate predictions of future SOC values to better plan optimal times to charge while ensuring that the SOC stays above the minimum SOC.

The network represents the flow of chargers from one task to another. Each vertex corresponds to a particular charger type either resting or charging a particular bus at a particular point in time. Integral to the concept of network flow is the idea that some quantity is moving (flowing) from one node (vertex) to another along the edges. The edges in this scenario denote the possible transitions for the chargers. When at a rest vertex, a charger may transition to a subsequent rest vertex or it may transition to begin charging a bus. When at a charging vertex, a charger may transition to a subsequent charge vertex of the same bus, or it may transition to a rest vertex at the subsequent time. Due to route schedules, the buses are not always available to be charged and vertices for charging a bus are only present when the bus is available. An edge between an adjacent pair of charging vertices is present when the bus receives charge,

i.e. no charge is given to a bus if the bus does not stay at the charger. A simplified example of such a graph is shown in Figure 33.

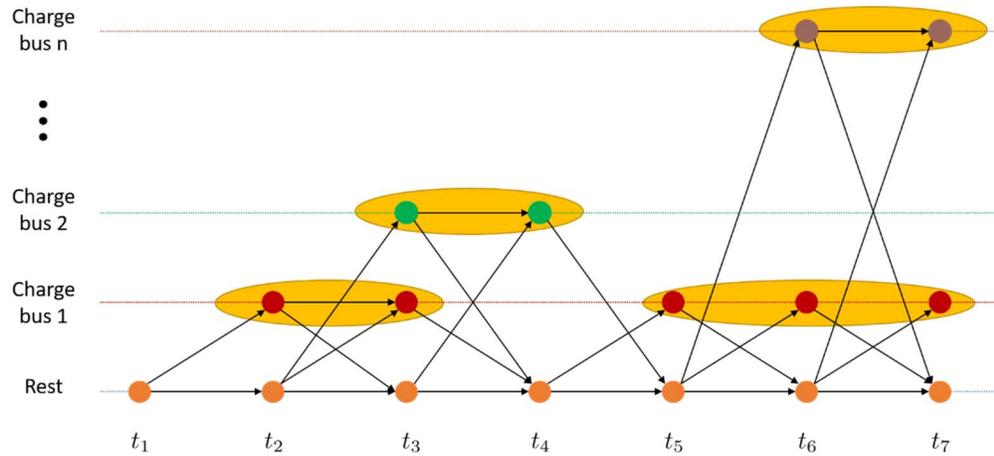


Figure 33 - A simple graph of the states and transitions for a single charger type. The bottom row corresponds to the charger at rest while the other rows each correspond to charging a particular bus. Yellow ovals are placed around charging windows.

The amount of flow in the graph is constrained to the number of chargers of the type represented in the graph. This automatically ensures that no buses are scheduled for the same charger at the same time. Each of the edges that corresponds to charging a bus is limited to one unit of flow to ensure that a bus is not assigned to more than one charger at a time. In addition, each charge window (the yellow ovals) is limited to one unit of flow, i.e., a bus can only charge one time while at the station. This prevents a bus charging for some time, leaving the charger, and then needing to drive back to charge again.

A network flow problem like this can be modeled with a mixed-integer linear program. An integer variable x_i represents the amount of flow through each edge and auxiliary variables track the bus charge level ($s_{j,k}$) and charge gain ($g_{j,k,l}$) for each time step, k . Constraints are added to ensure that the charge levels stay within operable ranges, the buses charge at most once per visit to the station, and the number of simultaneously charging buses does not exceed the number of available charging stations, amongst others. The demand and consumption aspects of the rate schedule are directly modeled in the objective function.

Conversion to this mixed-integer linear program allows the use of typical optimization software to generate solutions. We've chosen to utilize the commercial (free for academic use) Gurobi optimization software due to its convenient interface and high performance. It is worth noting that there is other software, including free, open-source projects, that can fill the same role.

For two buses, planning the whole day with a five-minute resolution can take around 30 seconds to solve. For thirty buses, however, planning the whole day to optimality can take much longer (ten minutes of solve time results in a solution that is within at least 20% of optimal, with a resolution of ten minutes). This means that while a two-bus problem can be solved fully and quickly enough to be usable in real time, unmodified, this approach does not scale well to more buses. By solving a problem with a shorter horizon, the solve time can be

reduced as much as needed (albeit at the cost of the ability to “plan ahead”). This allows the smart load management system to adapt in real time to variations in the bus schedules, charge rates, etc.

4.4 Field Evaluation

Two simulated scenarios are given for proof of concept. The first consists of a 30 bus scenario with six overhead fast chargers and a slow charger available for each bus in the depot. This scenario demonstrates the scalability of the proposed approach. The second scenario consists of two buses with two overhead chargers and two depot chargers to better model the scenario to be demonstrated.

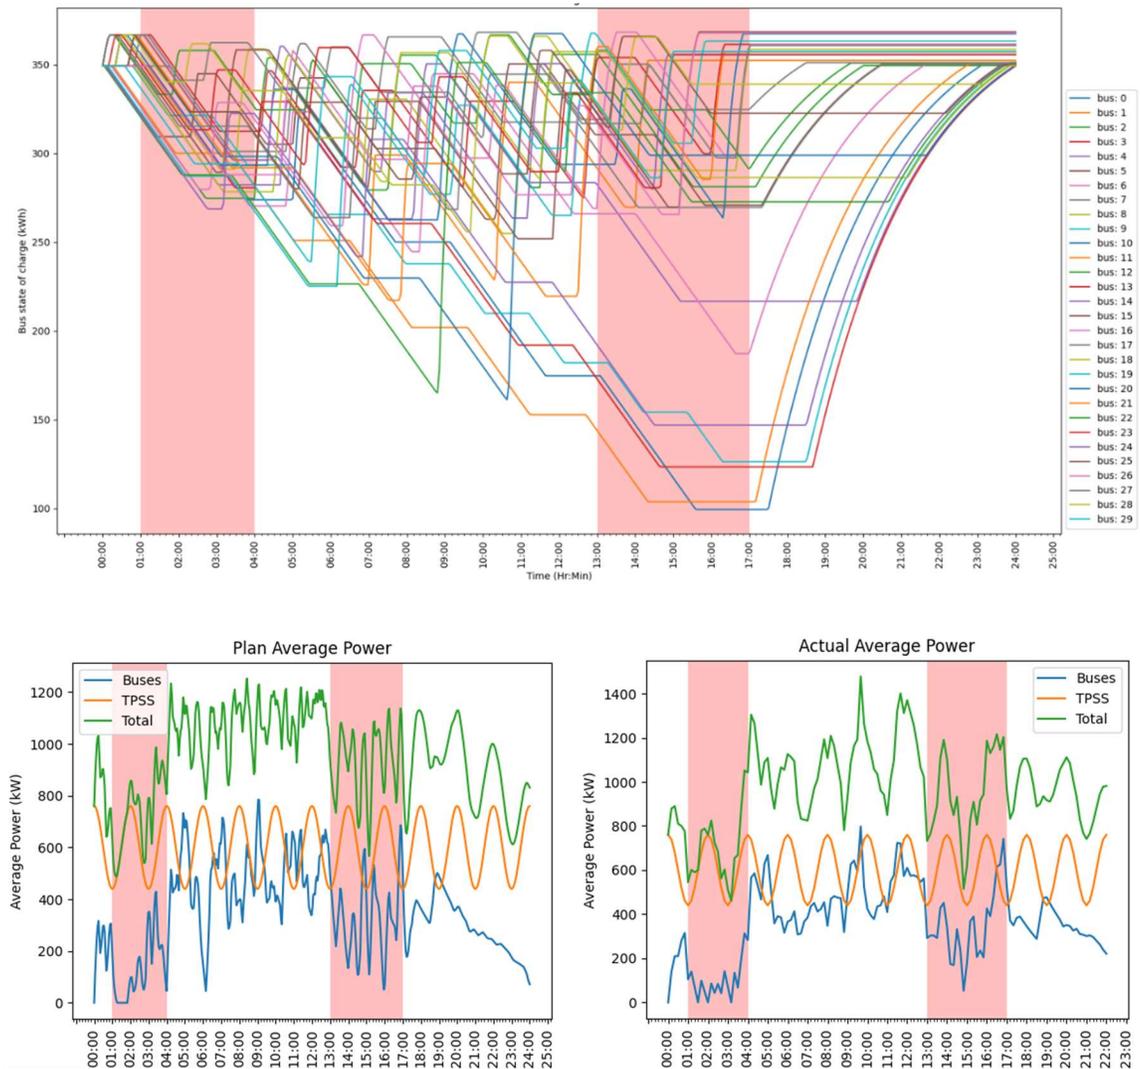


Figure 34 - The top image shows the charge for each of the 30 buses over a 24 hour period. The bottom left shows the power usage of the original 24 hour plan and the bottom right shows the resulting power usage after execution. In each of the figures, the pink zones correspond to peak charge times.

The **30 bus scenario** was produced using random schedules to represent the availability of each bus to be charged as well as a sinusoidal TPSS power usage signal. It was assumed that

each battery has a 388 kWh capacity. Constraints were included to ensure that each bus stay above 20% battery charge, 77 kWh, and return to 90% charge, 349 kWh, by the end of the 24 hour period. For battery health, the charge was not to exceed 95%, 369 kWh. The charge plan and execution can be seen in Figure 34. The top image shows the charge for each bus. It can be seen that each bus stays between 77 and 369 kWh of charge and returns to at least 349 kWh of charge by the day's end. The exponential charge rate can especially be observed at day's end when many buses are using the depot chargers.

The bottom images show planned and actual power usage. The bottom left image shows the planned power usage from the long horizon planner. Note that the planned power usage for the bus charging has peaks during the troughs of the TPSS usage, demonstrating a balance of power usage. The bus power usage is also lower during peak times. The bottom right image shows the actual simulated power usage. The receding horizon for these results consisted of a two hour period and solve time never exceeded three seconds. While the receding horizon planning does have similar attributes of avoiding peak windows and scheduling peak charge times during troughs of the TPSS usage, it does not do as well as the full 24 hour horizon planner. Future work will be investigating the results of problem formulation, planning resolution, and planning horizon times on the end performance of the receding horizon planner.

The **two bus scenario** was again produced using random schedules to represent the availability of each bus with the TPSS power usage and constraints being defined as before. Note that the full 24 hour horizon can be planned in less than 30 seconds. As the day progresses, the diminishing horizon planning time decreases. The results can be seen in Figure 35. Two versions of the two bus scenario were run for comparison. The first scenario did not take the TPSS information into account and the results are shown in the top and bottom left images of Figure 35. Note that without TPSS, there were only four uses of the fast chargers, leaving much of the charging to the depot charges at the end of the day. With the TPSS information considered, there are many more charging events throughout the day, with the bus charging power usage peaks occurring in the troughs of the TPSS power usage.

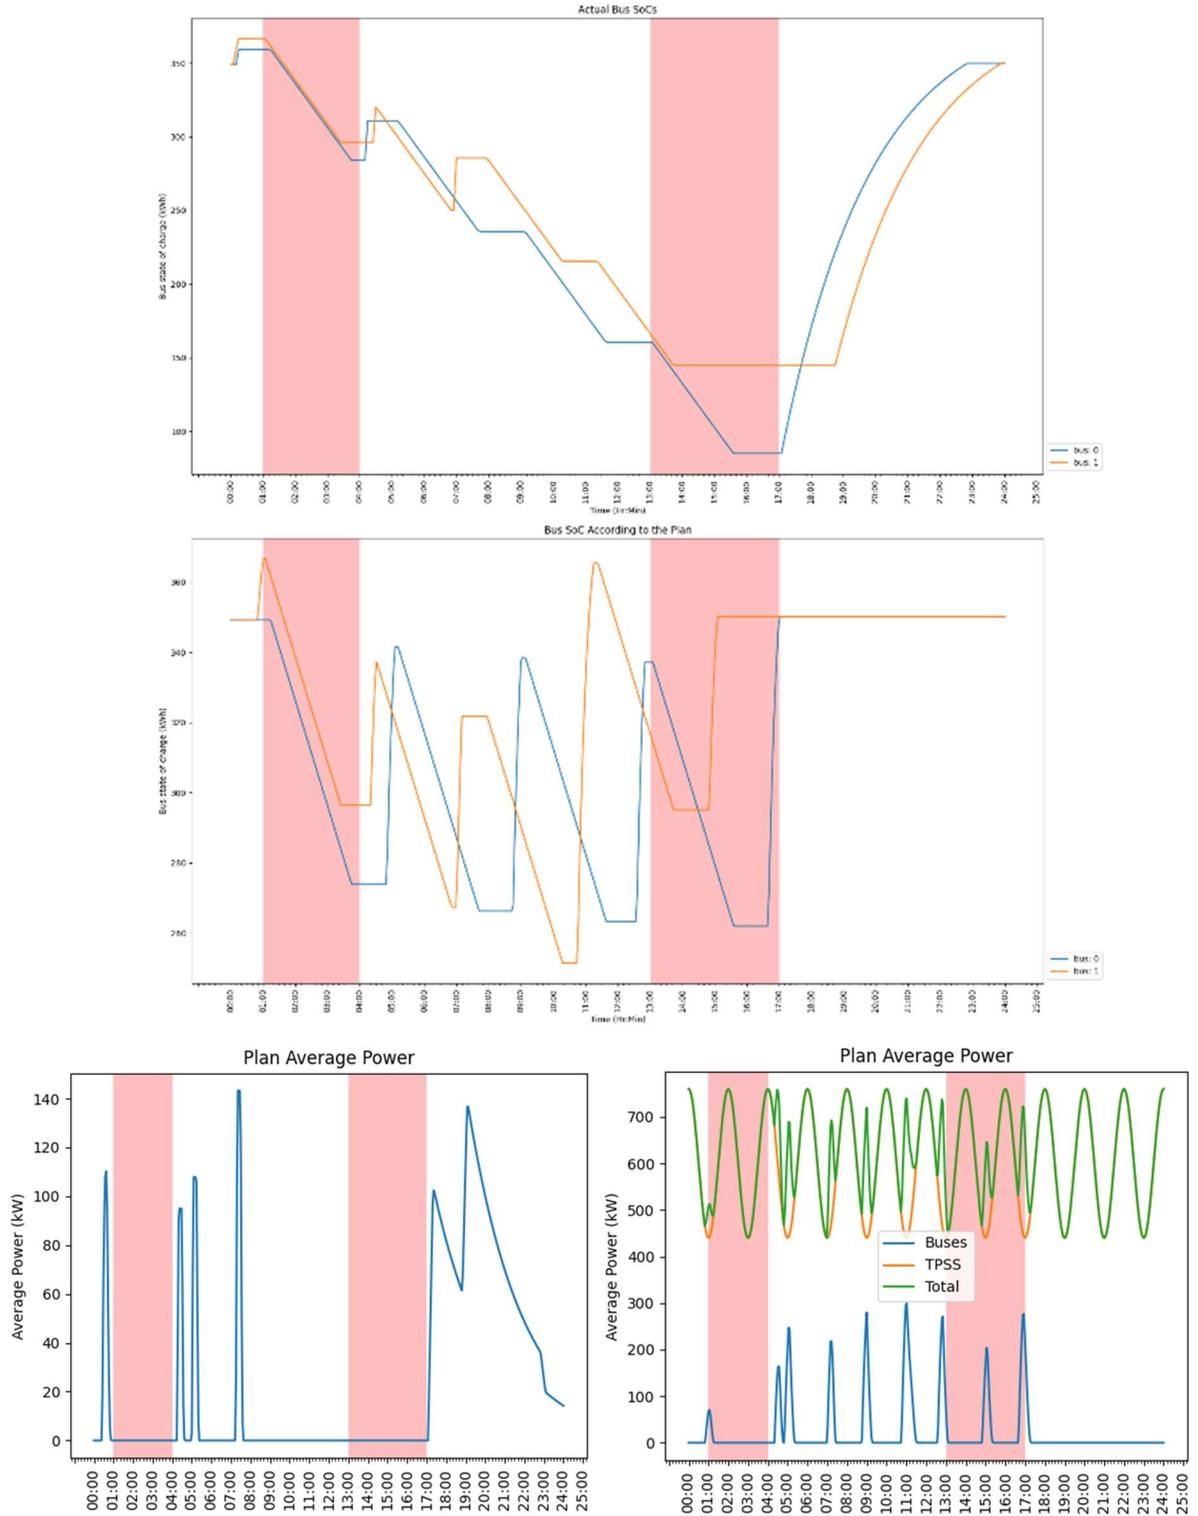


Figure 35 - Top image shows the planned charge for each bus without consideration of the TPSS. The middle image shows the planned charge for each bus while considering the TPSS. The bottom row shows the planned average power usage with the left being without TPSS

With the simulated results for optimized charging demonstrating improved power utilization

for the UTA site, USU took the next steps to coordinate deployment of the software. The demonstration required all active inputs for data at the UTA site as explained previously in this Task section. USU engineers were onsite at UTA installing cellular devices for data collection at the TPSS.



Figure 36 - Upgraded Red Lion device to TPSS Siemens system. USU connected a cellular device to pull data for algorithm input.

Further coordination occurred with the UTA mechanical support team at the bus depot, to understand the charging procedures that each bus operator follows when connecting to the overhead chargers and initiating a charge session. This was important so we could integrate our control process so as not to effect standard operations for the bus drivers. To ensure no-interruption or additional training requirements, bus operators were instructed to follow UTA's current SOPs and connect to the charger each time they arrived to the Intermodal Hub. The USU algorithm would then identify, through planning, if a charge should be administered and what power level would be provided.



Figure 37 - UTA electric bus at overhead charge station



Figure 38 - TRAX arriving at central station

Having established this communication and after securing required datasets for the deployment of the system architecture. We conducted testing at USU with the server control. Through these efforts we recognized additional challenges with command control over the commercial ABB chargers.

The ABB chargers maintain high-security standards; establishing connections to these devices from outside ABB required significant software development. Much of the task of initial connection involved ABB engineers working in tandem to help route charger data traffic to USU. Our servers have visibility into both UTA overhead chargers and USU local ABB fast-charge systems (UTA systems are marked in red – Figure 39). To ensure that best industry security standards were followed, a network subnet was created only accessible to a single server. All developed capabilities and connection software are hidden from external access.

microgrid	6u	IPv4	81098	0t0	TCP 129.123.202.150:9001 (LISTEN)
microgrid	7u	IPv4	78707	0t0	TCP 129.123.202.150:9001->46.108.235.154:28957 (ESTABLISHED)
microgrid	8u	IPv4	81102	0t0	TCP 129.123.202.150:9001->46.108.235.139:29220 (ESTABLISHED)
microgrid	9u	IPv4	81104	0t0	TCP 129.123.202.150:9001->46.108.235.239:11887 (ESTABLISHED)
microgrid	10u	IPv4	78705	0t0	TCP 129.123.202.150:9001->46.108.235.104:6044 (ESTABLISHED)

Figure 39 - Connections Established on Secure UTA Charger Devices

We recognized that this technology and communication protocol is in active development with ABB. Many features of OCPP are not yet integrated or implemented by ABB engineers, with USU developing many initial tests and use cases for the OCPP connection. ABB's recent hardware updates that updated communication protocols (3G to 4G) and internal software caused all previously established connections and software to fail.

USU has developed and continues to build several patches and documentation to enable more OCPP functionality on ABB charger systems. Connections have been re-established, and the availability status, charger name/ID, and some diagnostic data are now available (see Figure 41). New protocols do not easily accept prior working efforts in charge or curtailment commands. A new API and documentation from ABB systems is needed to be built to allow for new commands to be properly sent and received, as commands are sensitive and debugging code has required extensive effort.

USU has continued research in this area, and is active with the ABB research team to fully explore and establish a robust and controllable system. As this area of research is new there must be coordination with charge equipment OEM's before full deployment and

demonstration at a site can be achieved. We anticipate the continued development and validation of the demonstrated system-managed charge results through simulation.

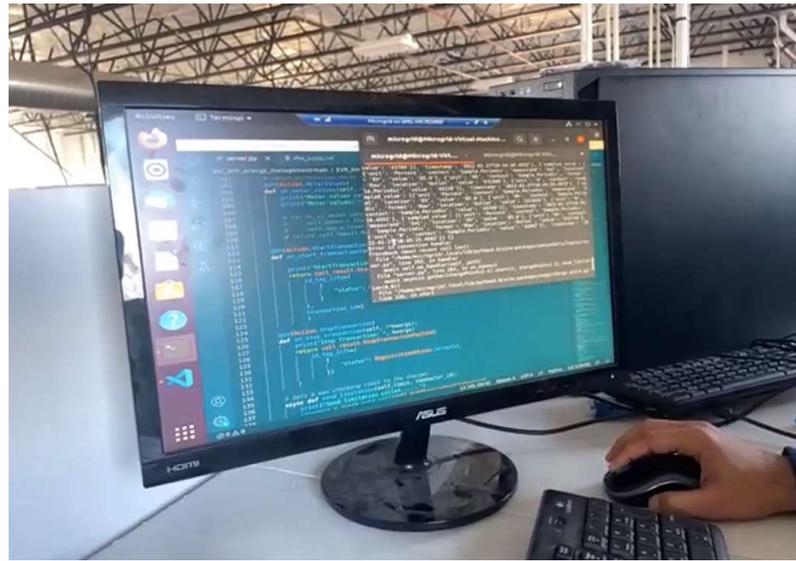


Figure 40 - OCPP server script running at USU

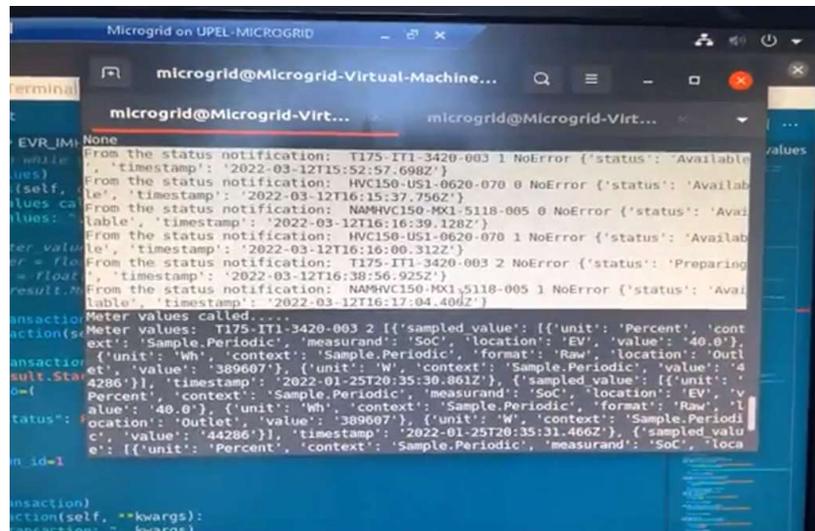


Figure 41 – Script verification that all chargers are connected to USU OCPP server. T175-IT1-003 and -016 are USU ABB chargers. NAMHVC150 and HVC150 are UTA overhead chargers.

5 Cost-benefit analysis

5.1 Cost-benefit analysis

Overview

A cost-benefit analysis has been performed that evaluates the impact of additive load scenarios on offsetting the cost of infrastructure to Rocky Mountain Power. These loads are in addition to what is currently in place at the UTA Intermodal Hub Site i.e., TRAX, a CNG compressor, two overhead bus chargers, and three ABB chargers in the UTA Central Garage bus depot, along with smaller loads such as the snowmelt and lighting fixtures.

The TRAX light rail experiences a high yet intermittent peak load that is a driver of demand

charges. Methods to offset these resultant demand charges include increasing the utilization of the supporting grid infrastructure by expanding facilities for the charging of electric buses, private vehicles, TNCs and local delivery fleets. These load scenarios, as given in Table 4, are treated as additive to present loads at the Hub. Installation of onsite solar power generation and battery storage is also explored parallelly to support load and mitigate demand charges. A prime consideration here is to determine revenue streams from these load scenarios and the business arrangements to support them.

Table 4 - Load scenarios as defined for the traditional cost-benefit analysis

Scenarios	Load Components
BAU	TRAX, CNG compressor, e-bus charging
Scenario 1	Expanded e-bus fleet
Scenario 2	Public Charging (L2, DCFC)
Scenario 3	TNC, local taxi, delivery fleet charging (DCFC)
Scenario 4	Onsite solar and battery storage

This analysis has achieved formulation of the cost-benefit model and application of the model for the Business-as-Usual (BAU) scenario covering the load components in operation at the UTA intermodal hub. Herein, we have collected and analyzed electricity data for existing loads, identified the costs of substation infrastructure upgrades over the next five years, and developed a methodology that takes into account not just future electric loads but also all potential revenue streams in order to inform a comprehensive rate-based payback approach.

The cost-benefit model developed under this project is currently being extended under the WSEV@Scale project to further quantify the electric loads attributed to Scenarios 1 to 3. This will further inform the economic feasibility of these additional vehicle charging scenarios. Parallelly, we are working to identify relevant stakeholders other than RMP and UTA and their interactions to determine possible business arrangements to support deployment of these scenarios.

Methodology

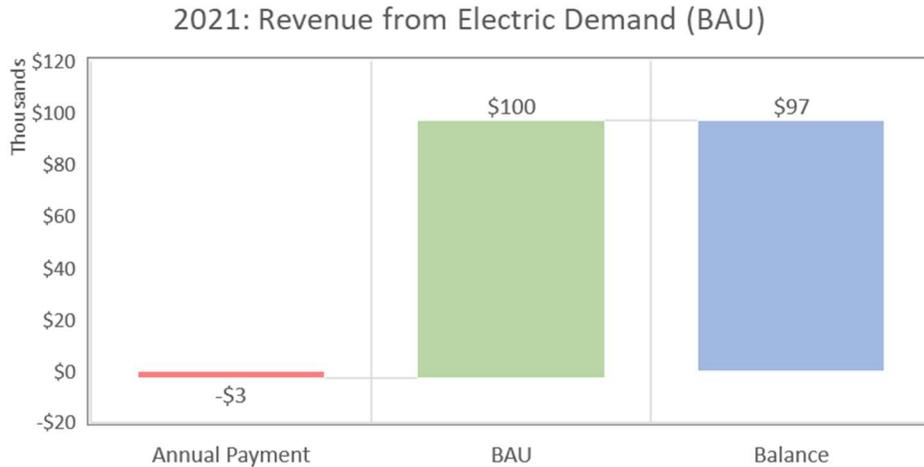
The cost-benefit model is set up with the objective to calculate payback from additional electric loads that increase utilization of the substation infrastructure upgrades at the Intermodal Hub. The results from this are to be used in developing revenue streams to recover costs of deploying these load scenarios through innovative business arrangements. The approach aims to inform creation of a rate-based payback system to cover infrastructure costs.

- We employ a Discounted Cash Flow (DCF) model to calculate annualized payments for infrastructure investments.
- To be consistent with RMP’s economic practices, we use the lifetime of infrastructure and Weighted Average Cost of Capital (WACC) as appropriate for the organization.
- These financial indicators are used to calculate the Net Present Value of the investment to be recovered in the given financial period.

- Infrastructure investments are accounted for annually and correspond to the incremental load at the Hub. The model normalizes these annualized payments over the corresponding annual incremental electric load, leading to unit cost metrics in \$ per kW and \$ per kWh.
- Revenue to the utility consists of the energy usage and subsequent billing of these load scenarios. Annual electricity usage revenue is calculated according to the existing RMP tariff that UTA is subject to i.e., Schedule 6 and 6A General Service Rate Cards.
- It is assumed that an additional charge may be applied to the billing of these loads at the Intermodal Hub to supplement revenue from electricity usage in covering costs. This would either build upon the service contract agreement currently in place or call for a new Master Electric Service Agreement between RMP and UTA. The unit cost metrics as described above can be used to estimate this additional charge in the service agreement, informing rate-design to recover infrastructure costs for the utility.

Results

For the BAU scenario, the model determines the \$ per kWh unit cost to RMP for substation infrastructure upgrades, given planned incremental load additions out to 2027 from RMP’s Interconnection impact study for a Load service Request by the UTA. According to the DCF model, substation infrastructure upgrades require an annualized payback of \$952,625 (amounting to \$0.95 per kWh) beginning in 2027, accounting for all planned investments.



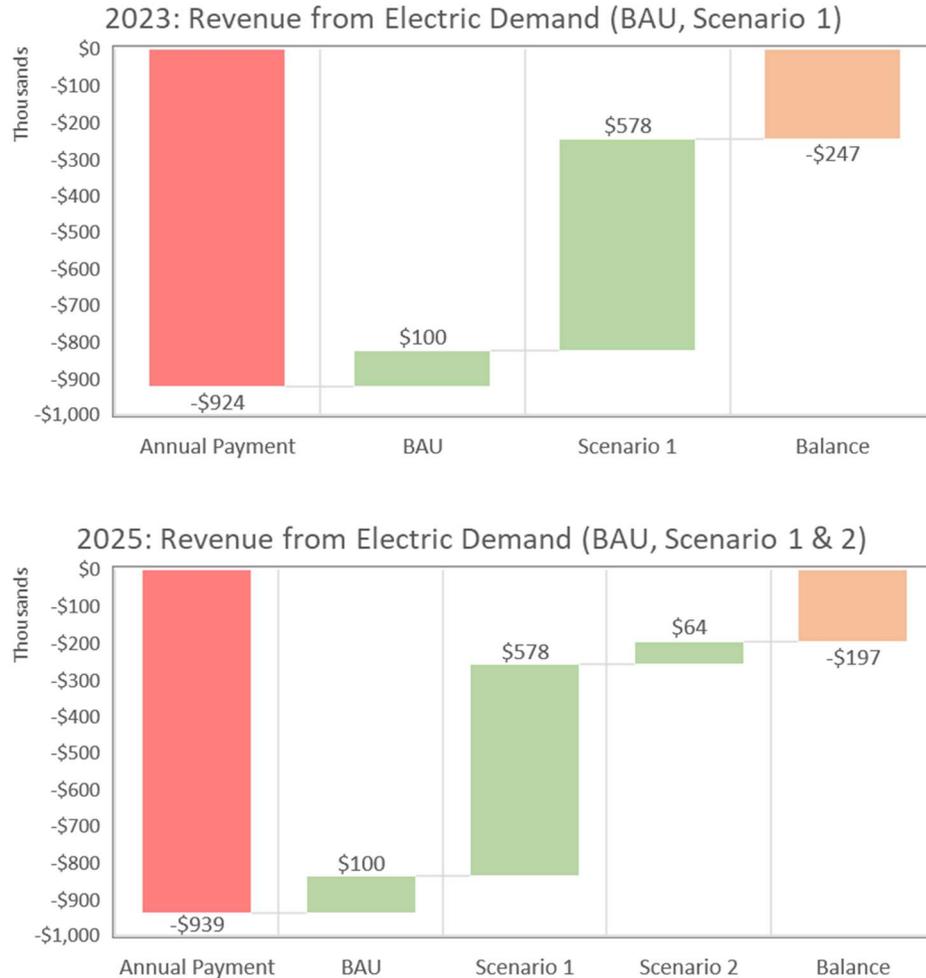


Figure 42 - The three charts above show expected impact of estimated annual revenue from electricity on offsetting annual payback on substation infrastructure investment to the utility. Scenario 1 and 2 correspond to the added electric load from expanded bus char

Electricity revenue calculations for present metered locations at the Intermodal Hub indicate that current annual electricity usage will not be sufficient to cover this cost through time Figure 42). Annual electricity billing from the three largest loads – TRAX, the Greyhound Station, and the overhead bus chargers – amounts to \$100,000 on average. Once simplified, the TRAX light rail billed on Schedule 6A and the Greyhound station on Schedule 6 offer revenue to RMP at a monthly average of \$0.03 to 0.07 per kWh, whereas the overhead bus chargers on Schedule 6 are billed at an average of \$1.01 per kWh. This can be attributed to the 270 kW average peak demand of the chargers and the demand-sensitive nature of the Schedule 6 tariff. Figure 42 shows the ability of electric revenue from loads at the Intermodal Hub to offset annualized investment costs. As a preliminary calculation, we assume Scenario 1 to include bus charging expanded to 10 overhead chargers, and Scenario 2 to include public charging with five 6.5 kW Level 2 and five 50 kW DC Fast Chargers installed at the hub. These additional loads help offset an additional 68% of the annual infrastructure investment cost, in addition to BAU. For the purpose of this phase of the task, we have assumed a 10% utilization for Level 2 chargers and 5% for DCFC. It is anticipated that the load components of Scenarios

2 and 3 will contribute to a larger electricity usage than estimated at this stage, further increasing electricity revenue and in turn offsetting a larger share of the annualized cost. It is a preliminary conclusion that the electricity usage of high-power electric vehicle charging infrastructure as to be developed for Scenarios 2 and 3 is expected to have a positive impact on this cost-benefit analysis.

6 Assessment of the potential for future deployment

The Intermodal Hub Project has demonstrated the ability to safely and reliably operate a power balance and demand response program that both increases revenue from utility infrastructure investments and reduces energy costs for utility customers that would otherwise have high peak demand. This has been shown for the case of sharing the loads from multiple modes of transit with complementary schedules and priorities. As has been demonstrated, power balance was achieved without compromising customer needs. The TRAX line peak demand took highest priority, while its short intervals had minimal impact on balancing the loads with overhead bus charging and the associated bus schedules.

As was shown, the business as usual (BAU) scenario results in very poor utilization and high cost for infrastructure and generation. This high cost is passed on to the utility customers through a combination of upfront upgrade costs, minimum service charges, and peak demand charges. These cost barriers are expected to significantly limit or at least slow EV adoption and private investment in charging infrastructure.

The cost benefit analysis has shown a very significant advantage in Scenario 1 by leveraging and balancing the load with the combined loads of the TRAX line and an expanded e-bus fleet. The results of the Intermodal Hub Project will directly extend to further expansion at this site. UTA is anticipated to expand their existing bus fleet to 50 electric buses at the central station site. By applying the planning and operation tools from the project, UTA will be able to reduce the extent of utility upgrades required to support the bus fleet and reduce operating costs beyond those shown under Scenario 1 in this study due to even greater utilization of the infrastructure.

Furthermore, initial results show similar advantages in Scenarios 2 and 3, and application to broader sites. By adding public L2 and DC fast charging, and by providing incentives for coordinated power balance and demand response from TNC, local taxi, and local fleets, including delivery trucks, state, local and emergency services, even higher levels of utilization can be achieved. Although further studies are needed to demonstrate integration at Scenarios 2 and 3, the initial results clearly show the advantages of power balance and load management across modes of transit and types of users. In addition, although our initial analysis for public charging assumed typical 10% utilization for L2 and 5% for DC fast charging, it is anticipated that significantly higher levels of utilization from public charging can be achieved in the future with potential concepts such as incentives, price signals, or distribution level futures contracts, and by combining public charging with TNC, local taxi, and other fleet vehicle charging.

The key site characteristics for application of the project solutions are existing or planned loads with high peak to average power ratio and nearby loads with relatively consistent energy needs and some flexibility. As one important example, these characteristics are common of utility customers considering the transition to electric fleets. In these cases, initial costs for supporting the charging needs of a relatively small fleet are very high for a single customer. By supporting solutions based on the findings of the Intermodal Hub Project, the utility can help encourage the transition to electric for customer fleets and allow neighboring customers to balance loads and share infrastructure across multiple fleets and with public, transit, and TNC and taxi charging needs in the region. The increase in base load and

improved utilization of utility infrastructure helps the utility keep rates low for rate payers.

The market for electric vehicles in Utah is expected to grow significantly over the next two decades. Estimates of electric vehicles (EVs) by 2040 range from 200,000 without utility programs supporting infrastructure to well over 1 million with the utility providing early incentives and cost effective solutions to support charging infrastructure and low utility rates. The solutions developed under the Intermodal Hub Project could help drive accelerated adoption toward the high adoption curve estimates, where annual charging demand is expected to exceed 5 TWh by 2040 (more than 12% increase over existing load in Utah).

A direct example of opportunities for future deployment of the project results is to apply the technologies at additional UTA sites all along the TRAX line in the Salt Lake City area, as shown in Figure 43(a). The shared use of infrastructure for the TRAX line throughout the urban region can provide power balancing in support of additional on-route charging of UTA buses, electrification of UTA's micro-transit solutions, and public charging for TNC, taxi, local fleets, and private passenger vehicles. The concept can be further expanded to future electrification of FrontRunner along the I-15 corridor, as shown in Figure 43(b). With this expansion, the significantly higher power demands of a battery-electric heavy rail system would be cost effectively balanced with electric buses, micro-transit, TNC, taxi, and local fleets, and importantly, with expansion of electrified passenger and freight traffic all along I-15.

Finally, the concepts developed and piloted under the Intermodal Hub Project for power balancing and demand response of multi-modal transportation are expected to be of value to utility customers beyond transportation. With additional algorithm development, the solutions can be applied to buildings and building loads throughout connected communities and combined with electrified transportation solutions in the region. Once again, the key value is the ability to aggregate and co-manage loads with different characteristics and priority levels to reduce and defer infrastructure upgrade requirements and improve overall infrastructure utilization.

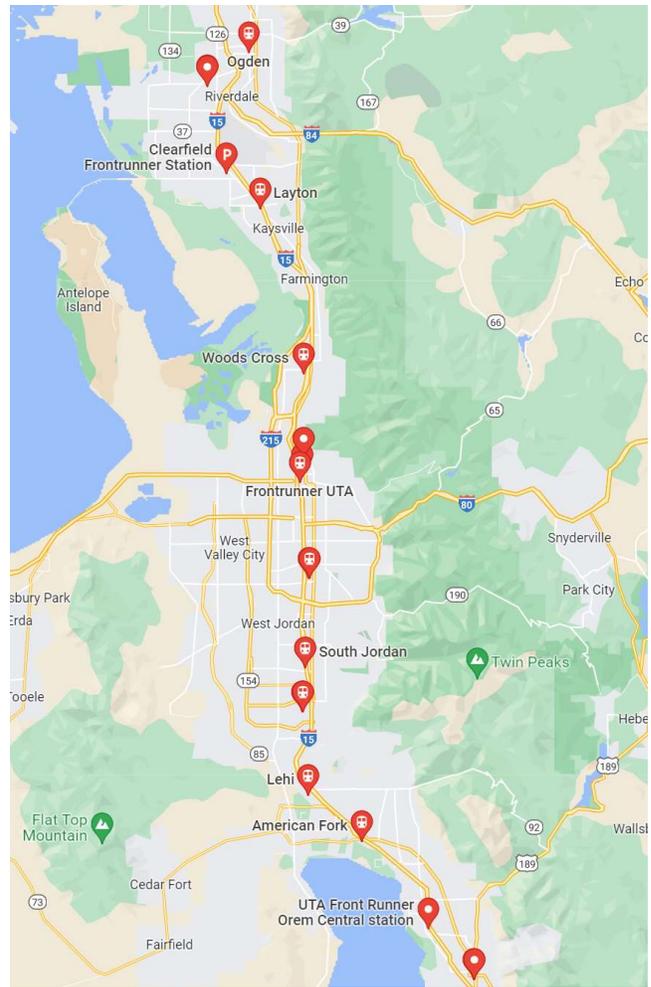
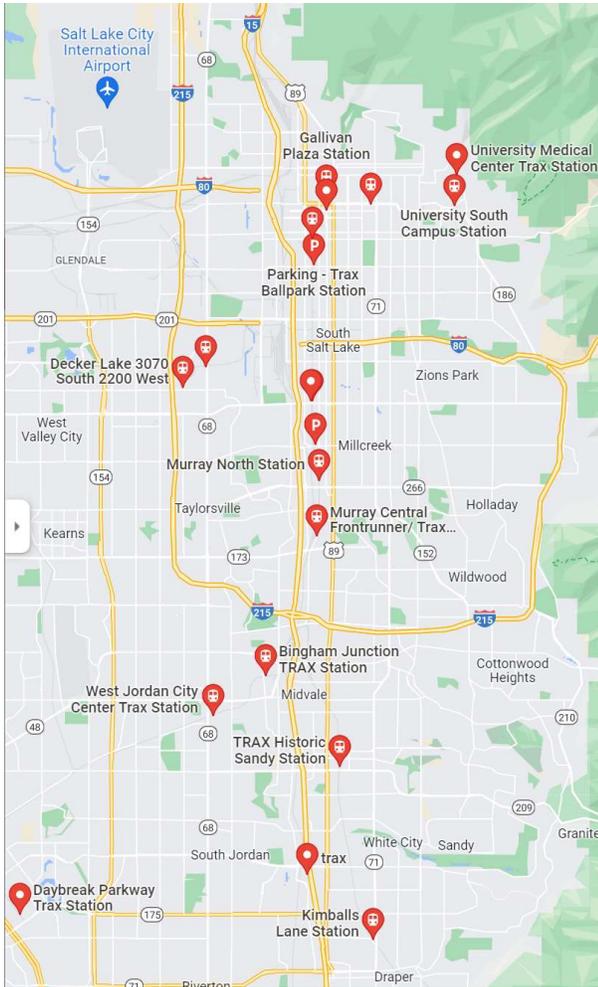


Figure 43 - Example opportunities for future expansion of the Intermodal Hub Project technologies with multi-modal support of rail, bus, and private and fleet vehicles. (a) Expansion throughout Salt Lake City at additional TRAX line stations in support of local urban traffic and (b) future batter-electric train replacement of FrontRunner along the I-15 corridor in support of broader passenger and freight electrification in the state

