Scheduling a Multi-Cable Electric Vehicle Charging Facility

Tony T. Tran[1], Mustafa K. Dogru[2], Ulas Ozen[2] and J. Christopher Beck[1]

[1]Department of Mechanical and Industrial Engineering
University of Toronto
{tran,jcb}@mie.utoronto.ca

[2]Alcatel-Lucent Bell Labs
{Mustafa.Dogru,Ulas.Ozen}@alcatel-lucent.com

Abstract

We consider scheduling electric vehicles in a charging facility where customers arrive dynamically and tend to park longer than their charge time. In this setting, it is reasonable and technologically feasible to have charging docks with multiple cables, although such docks do not currently exist in practice. Assuming such a dock design, we study three information conditions: we know the number of electric vehicles at each dock, we know stochastic information about arrival and charging requirements, and we are able to observe exact charging requirements for vehicles in the system. We formulate a continuous-time Markov decision process (CTMDP) to optimize the system performance under the first two conditions and demonstrate that it does not scale to realistic-size problems with multiple docks. However, a single-dock version of the CTMDP is tractable. We propose and numerically evaluate a number of admission and scheduling schemes building on both the single-dock CTMDP and approaches from the scheduling literature under each of the three information conditions. Our results demonstrate (i) the value of a multi-cable dock, (ii) the importance of obtaining actual charging requirement information, and (iii) the integral role of admission and scheduling policies based on available information to improve performance.

1 Introduction

Advances in battery, electric engine and charging technologies have resulted in significant improvements in the performance of electric vehicles (EVs) in terms of range, charging time, etc. Although reduced emissions and lower fuel and maintenance costs over their lifetime favor EV adoption over internal combustion engine vehicles (CVs), range anxiety prevents more people from owning an EV. Range anxiety is the fear of being stranded because an EV has insufficient capacity to reach a destination (Tate, Harpster, and Savagian 2008). Unlike CVs which run on gasoline, an EV requires an electrical power source to recharge its battery and completely recharging a fully depleted battery can take from half an hour to almost a full day. For example, a Nissan Leaf with a 24 kWh capacity battery has a range of about 73 miles on a single charge and requires 16-18 hours to fully charge from a depleted battery under level I AC (120V) chargers. Level II AC (240V) chargers decrease the required charge time to 7 hours and level III DC (500+V) chargers further reduces the time to approximately half an hour.[1]

To address range anxiety, charging stations are being placed in convenient locations including highway rest stops and gas stations. It is also becoming popular to place charging stations in parking lots. Cars generally spend a large amount of time in parking lots, whether it is a shopping mall, an airport, or work place. These charging stations provide a convenient way to charge a battery by integrating the charging into time periods which drivers are naturally occupied.

Current charging docks have a single cable and can be connected to one EV. In a gas station or rest stop, one would expect customers to leave when charged and so a new car can be connected immediately. However, in a parking lot, a vehicle may be connected to a cable well after charging has been completed. A charging dock which incorporates multiple cables will allow as many connected cars as there are cables, even if only one car can be charged at a time. With a multi-cable dock, a car may complete its charge and stay connected, while another EV immediately begins charging. Such a dock is an economical way to improve effective charging capacity without purchasing more docks. For example, the annual cost (purchase + maintenance) of a level II AC charger ranges between \$900 and \$5000 USD over a 10-year life cycle, and only about 20% of the total cost is due to initial capital investment (Botsford 2012). Although the cost of adding a cable to a dock is not negligible, we expect much lower maintenance cost for the multi-cable dock design in comparison to an equivalent system with multiple single-cable docks. Hence, the multi-cable dock design decreases the total cost of a charging facility due to fewer docks needed to purchase and maintain, and lower initial installation cost.

In this paper, assuming multi-cable docks, we study the admission and scheduling problem associated with management of an EV charging facility. EVs arrive dynamically over time and can be plugged into an available cable to be charged. Admission and scheduling decisions must be made immediately upon arrival of an EV and the system manager aims to minimize the costs associated with rejecting and delaying customers. Given the relatively new application of

[1]http://www.nissan.ca/vehicles/ms/leaf/en.

EV charging facilities and EVs themselves, available functionality varies. In particular, the information available to a system manager from both his/her docks and the customers' EVs will vary. We therefore propose three information availability characteristics: (i) the number of EVs at each dock is known, (ii) stochastic information is available about EV arrival and charging rates, and (iii) exact charging requirements for all EVs in and arriving to the system are known. We create and study policies for admission and scheduling in each information environment and compare the performance of a multi-cable charging dock facility.

Our study demonstrates:

- the value of a multi-cable dock for parking lot facilities,
- the importance of obtaining actual charging requirement information from EVs, and
- the increase in system performance arising from intelligent admission and scheduling policies.

In the following section, the charging facility we study is described in detail. Section 3 presents a continuous-time Markov decision process (CTMDP) for the system. However, the model suffers from the curse of dimensionality and hence does not scale well to real life problems. Thus, heuristic methods for admission and scheduling decisions are introduced in Section 4. Experimental results are presented in Section 5, followed by a discussion in Section 6. Some related work on EV charging can be found in Section 7 and Section 8 concludes the paper.

2 System Model

We consider an EV charging facility with $N \in \mathbb{N}$ docks, each with $K \in \mathbb{N}$ cables. A cable connects a dock to a car and enables charging. However, being connected does not mean that the car is able to immediately start charging. Each dock is limited to charging a single car at a time.

The parking lot system studied assumes cars arrive dynamically following a Poisson process with rate λ. The amount of charging time each EV requires is exponentially distributed with mean μ^{-1}. In order for the vehicle to leave the system, two conditions must be met: 1) the required charge is completed and 2) the deadline specified by the driver is reached. We assume the deadline is exactly L time units after the arrival of the EV and represents the time at which the customer has agreed to return to remove the EV. This is a simplification of the real system which one can think of as having customers with different deadlines. However, our assumption represents a parking lot that sells an exact amount of parking time to all customers, but each customer will have different charging times. If a vehicle is charged before the deadline, it must wait until the deadline before it can exit because, typically, the driver will not return for the EV before the deadline. On the contrary, if charge completion occurs after the deadline, the EV is delayed and must wait until the charge completes before exiting the system.[2]

[2] An alternative system could have EVs leaving at the time of the deadline regardless of charge. The models presented in this paper can just as easily handle these systems with minor alterations.

We assume three information conditions for our system. The first is referred to as the *cardinality* condition: it is known how many EVs are at each dock and of those, whether or not a vehicle is delayed or charged. Under the second, *stochastic*, condition, the arrival (λ) and charging (μ) rates and their distributions are available. It corresponds to assumptions common to stochastic modelling (Puterman 1994) and queueing theory (Gross and Harris 1998). Finally, we wish to consider information natural to the scheduling community (Pinedo 2008) which tends to include deterministic information about job durations and, often, arrival times. While deterministic arrival times are unrealistic in our application, it is reasonable to assume that charging time information is known upon an EV arrival. For example, the charging time can be found from either requesting the customer give the charge level they wish to purchase or by having a wireless transmitter from the vehicle broadcast this information.[3] Therefore, our third condition, which we term *observable*, assumes that the actual remaining charging times of every EV in the system at each time point can be observed. For an arriving EV j, the charging time p_j is available upon arrival.

The system manager makes two decisions. The first is whether to accept or reject an incoming vehicle. If rejected, then there is a finite cost $c_r \geq 0$ for losing a customer. The second decision is how to schedule an accepted EV. Scheduling comprises of the decision of assigning a dock for an EV and determining the order that EVs are charged. If accepted, an EV is immediately assigned to an available cable and cannot be switched. When the owner returns to pick up his/her EV, if charging is not yet complete, the delay is penalized. If T_j is the tardiness of a late EV j, that is, the time between the EVs deadline and when its charge is completed, then the delay cost is $c_d T_j$ where c_d is finite and non-negative.

The system manager wants to find an admission and scheduling policy to minimize the overall system cost. However, the control a system manager has over a parking lot may vary. One can see in most common parking lots, customers arrive and choose a spot themselves. Here, a system manager would have no direct control over customers. Thus, an indirect method to control the system is by limiting the available spots (docks and cables). Although we do not explore the capacity planning problem, we will observe some of the effects of adding cables and docks to the system in our experiments. A system with moderate admission control could be seen as having a gate at the entrance of the parking facility to turn customers away. Once admitted, the customer is free to choose whichever spot they please. Finally, we can imagine a facility where customers must purchase a spot first and will then be assigned to a specific location. In this way, complete control over the admission and scheduling of a vehicle is possible upon arrival. Specific admission and scheduling policies will be discussed in Section 4.

[3] Such transmitters are already available (Botsford 2012), but not used widely.

3 Continuous-Time Markov Decision Process

We present a CTMDP model to handle the admission and scheduling of a charging facility when only cardinality and stochastic information is available. Our current definition of deadlines being a fixed L time units after arrival does not adhere to the memoryless requirement of a CTMDP. Therefore, we assume that deadlines are not deterministic, but exponentially distributed with mean L.[4] We further simplify the CTMDP representation by enforcing first-come, first-served (FCFS) ordering of EVs once assigned to a dock.

The state of the system at time t is represented by, $\mathbf{S}(t) = \{\mathbf{Q}(t), \mathbf{W}(t), \mathbf{D}(t)\}$. Here, $\mathbf{Q}(t)$, $\mathbf{W}(t)$, and $\mathbf{D}(t)$ are vectors of size N. $\mathbf{Q}(t)$ indicates the number of EVs in the system at time t that are waiting for a charge and not yet due on each of the N docks. $\mathbf{W}(t)$ represents the number of vehicles that have completed their charge, but are waiting for the deadline and $\mathbf{D}(t)$ is the number of vehicles that are not yet charged but have already reached their deadline, on each of the N docks. We represent the element in each of the vectors using a lowercase letter with index n to indicate the dock (e.g., the nth dock is fully described by $q_n(t)$, $w_n(t)$, and $d_n(t)$).

There are $N+1$ possible actions when an EV arrives. An action, $a \in A$, can either assign the EV to one of the N docks or to reject the vehicle. Therefore, $A \in \{0, 1, \ldots, N\}$ where 0 represents rejection. An EV cannot be assigned to a dock with no available cables (i.e., if $q_n(t) + w_n(t) + d_n(t) = K$). The cost function, $C(\mathbf{S}(t), a)$ defines the expected cost associated with action a in state $\mathbf{S}(t)$. When a vehicle is rejected, independent of the current state, the cost is c_r. If a vehicle is admitted, then we must calculate the expected cost associated with the additional vehicle for each particular state. We denote the time that an EV j completes its charge as ϕ_j. Thus, the delay cost of a vehicle is $\max\{0, (\phi_j - L)c_d\}$.

We calculate the expected delay of an accepted vehicle by conditioning on the state of the system at time t and the dock n that will be assigned the vehicle. Since there are $q_n(t)+d_n(t)$ vehicles not yet charged on dock n, admission of a new vehicle requires a total of $B = q_n(t) + d_n(t) + 1$ exponentially distributed charges until the arriving vehicle has completed its charge. B is the number of EVs present in the system that requires a charge plus the new arriving job. Thus, the expected delay given a state i and assignment to dock n is,

$$\mathrm{E}[\mathrm{delay}|\mathbf{S}(t), a = n] = \int_L^\infty (x - L) f(x; B, \mu) \mathrm{d}x,$$

where $f(x; B, \mu)$ is the density function of the Erlang distribution. This yields

$$\mathrm{E}[\mathrm{delay}|\mathbf{S}(t), a = n] = \frac{\left[\mu^{-1}\Gamma(B+1, \mu L) - L\Gamma(B, \mu L)\right]}{\Gamma(B)}.$$

[4]We found numerically through simulation that a system with deterministic deadlines does not behave differently from the calculated CTMDP with exponentially distributed deadlines under FCFS. Due to space restrictions, we do not present these details.

Here, $\Gamma(b)$ is the gamma function and $\Gamma(b, \mu L)$ is the upper incomplete gamma function. Therefore, for any action a, we know the expected delay, which we multiply by c_d to obtain the expected delay cost.

The transition rates depend on the current state of the system, $\{\mathbf{Q}(t), \mathbf{W}(t), \mathbf{D}(t)\}$. Transitions occur because of three types of events: EV arrival, charge completion, and meeting a deadline. Actions only affect transition rates for the arrival events; the other events are independent of the actions taken.

We define an N-sized vector $\mathbf{e}_n$ that has 1 as the nth element and the rest 0. A deadline can occur on any dock which has $q_n(t) + w_n(t) > 0$. If $w_n(t) > 0$, then a transition occurs with rate $\frac{q_n(t)+w_n(t)}{L}$ and will change the state to $\{\mathbf{Q}(t), \mathbf{W}(t) - \mathbf{e}_n, \mathbf{D}(t)\}$. If $w_n(t) = 0$, a transition occurs with rate $\frac{q_n(t)}{L}$ to state $\{\mathbf{Q}(t) - \mathbf{e}_n, \mathbf{W}(t), \mathbf{D}(t) + \mathbf{e}_n\}$. Charge completions can occur on any dock which has $q_n(t) + d_n(t) > 0$. If $q_n(t) + d_n(t) > 0$, a transition occurs with rate μ to $\{\mathbf{Q}(t), \mathbf{W}(t), \mathbf{D}(t) - \mathbf{e}_n\}$ if $d_n(t) > 0$, and $\{\mathbf{Q}(t) - \mathbf{e}_n, \mathbf{W}(t) + \mathbf{e}_n, \mathbf{D}(t)\}$ otherwise. For an arrival event, we must consider the action taken. If an EV is rejected, then there is no transition. If we decide to assign an arriving vehicle to the nth dock, then there is a transition rate of λ to $\{\mathbf{Q}(t) + \mathbf{e}_n, \mathbf{W}(t), \mathbf{D}(t)\}$. Since we consider exponentially distributed inter-arrival times, charging times, and deadlines, the system is memoryless and we can restrict the decision epochs to only the times when the state changes.

A policy, π, specifies the action, $a^\pi(\mathbf{S})$, for each state $\mathbf{S} = \{\mathbf{Q}, \mathbf{W}, \mathbf{D}\}$. We can use uniformization (Lippman 1975) to discretize the MDP and solve for an optimal policy using policy iteration (Howard 1960) since we have a finite state space with bounded costs (Puterman 1994).

The CTMDP suffers from the curse of dimensionality: solving the CTMDP for real life problems is intractable as the number of states grows exponentially. The number of states for any particular system is $(K+1)^{2N}$. With five cables and five docks, we see that there are more than 60 million states. Thus, such a model is intractable for parking facilities of even moderate capacity. Nevertheless, this model can guide us to heuristics that use stochastic information which we present in the following section.

4 Admission and Scheduling

We propose admission and scheduling policies to manage the charging facility for each of our information availabilities. Depending on the particular conditions of information availability, some policies may not be possible to perform.

4.1 Admission Policies

The admission policy decides whether to accept or reject an EV upon arrival. The policies consider each dock and decides which docks are able to be assigned the EV. We present three policies which represent systems that, respectively, use cardinality, stochastic, and observable information:

- Free Cable - A vehicle is admitted if there are available cables - i.e., if $\exists n \ : q_n(t) + w_n(t) + d_n(t) < K$. Any dock with an available cable may be assigned the EV.

- CTMDP1 - Consider a single-dock version of the model. Solve for the optimal single-dock policy using CTMDP of Section 3 with the same parameters of the original multi-dock model (μ, L, c_d, c_r) except an arrival rate of $\frac{\lambda}{N}$. Given the state of the system $\mathbf{S}$, for every dock $n = 1, 2, \ldots, N$, check whether an arriving EV would be admitted in state ($q_n(t)$,$w_n(t)$,$d_n(t)$) under the optimal single-dock policy. The docks that accept an arriving EV are the only ones that can be assigned the EV.
- Myopic - Using the charging times, calculate the delay cost of scheduling an EV on each dock. If the cost of accepting the EV on the dock is less than the cost of rejecting the EV, then accept and assign to one of these docks.

Although the admission policy will limit how one can assign an EV, it does not assign a dock.

4.2 Scheduling Policies

Once an EV is admitted, a policy is used to assign a dock. Again, each policy represents systems that, respectively, use cardinality, stochastic, and observable information.

- Random - Randomly choose among one of the possible docks determined by the admission policy.
- CTMDP2 - Similar to CTMDP1, restrict the CTMDP model to a single dock and solve the Bellman equations to find the expected cost of being in each state. From the set of possible docks as defined by the admission policy, choose the dock in a state that yields the minimum expected cost.
- Earliest - From the set of possible docks defined by the admission policy, choose the dock that will result in the earliest completion time for the EV if all other already assigned EVs complete charge first.

These policies represent different levels of control from no involvement, where we expect customers to enter and choose a cable randomly, to complete control where a customer is sent to a particular dock in order to maximize performance. Once assigned to a dock, EVs are charged in FCFS order.

A system manager couples an admission policy with a scheduling policy to control the charging facility. It is obvious that the information availability would limit his/her choice of policies above. For example, CTMDP1-Earliest can only be used if cardinality, stochastic, and observable information conditions are met.

5 Experimental Results

We simulate the charging facility to observe the effects of using multiple cables and different policy combinations. For each experiment, 10 instances of 100,000 time units are simulated for every admission-scheduling policy pair. In all experiments, customers set a deadline of exactly 1 time unit after their arrival ($L = 1$) and docks have a charging rate of $\mu = 6$. For example, if our time unit is 3 hours, the parameters represents a parking lot which customers park for 3 hours and request on average 30 minutes charging time.

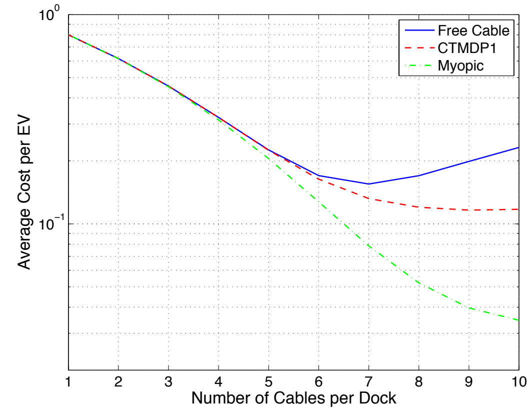

Figure 1: Experiment 1 - Single Dock Charging Facility.

As mentioned earlier, we wish to understand the effects of information availability. Since the underlying system does not change, only the available information, we can observe how having certain information affects performance.

The first experiment is a single dock charging facility. EVs arrive to this system with a rate $\lambda = 4$ and K varies from 1 to 10. The cost of rejecting an EV is $c_r = 1$ and the delay cost is $c_d = 5$. Since there are no scheduling decisions to be made, this system only tests the admission policies, hence, comparing the three information availabilities.

Figure 1 presents simulation results for the single dock system. In general, we see a large decrease in cost with additional cables due to a significant increase in accepted EVs. The cost gap between CTMDP1 and Myopic, especially for $K > 6$, shows the extent of performance improvement that can be achieved by obtaining the actual charging time information.

Interestingly, the average cost per EV of the Free Cable policy increases with eight or more cables. To get a better idea of why the increase in cost occurs, we can think of accepting an EV when there are 7 other EVs waiting for a charge. In this scenario, it is likely that completing eight charges will require more time than the deadline allows because the docks are expected to only charge 6 EVs in 1 time period. Therefore, rejecting is generally a better choice. The Free Cable policy does not do so and continues to accept EVs when there are free cables.

The second experiment looks at a multi-cable, multiple dock system. In this system, there are ten identical docks with between one and ten cables each. Vehicles arrive at a rate of $\lambda = 50$. The rejection cost is $c_r = 1$ and the delay cost is $c_d = 5$. Results are shown in Figure 2 for each combination of admission and scheduling policy. Note that CTMDP1 does not always perform better than Free Cable. With observable information, using the Earliest policy favours Free Cable. Figure 2 illustrates the importance of information availability. The strong performance of Myopic-Earliest shows the clear advantages of having observable information. Further, obtaining some control of the system is

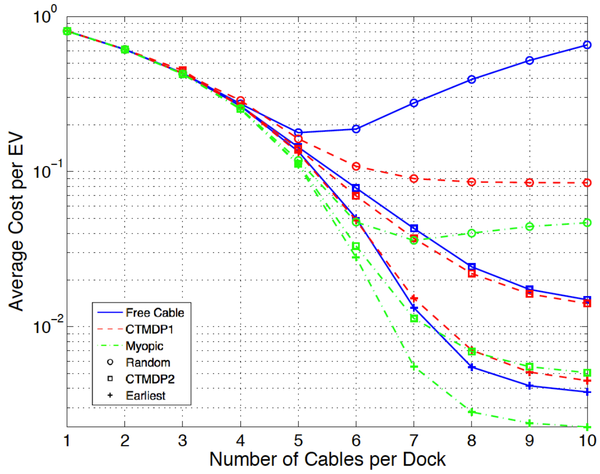

Figure 2: Experiment 2 - Multiple Dock Charging Facility.

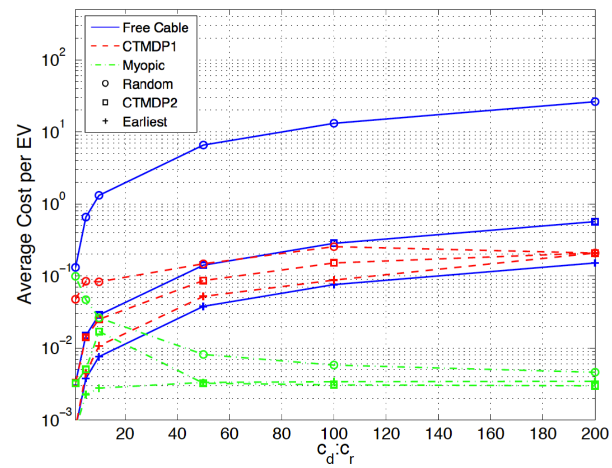

Figure 3: Experiment 3 - Multiple Dock Charging Facility.

quite important as Free Cable-Random is found to perform very poorly once there are seven or more cables. In fact, even Myopic-Random suffers when K increases since there is less control over the scheduling of EVs.

To further study the effects of the system parameters, we experiment with varying the cost structure. Using the parameters from the multiple dock facility of the previous experiment, we fix the number of cables to ten and vary c_d between 1 and 200. Figure 3 shows the results of this experiment.

We see for most policy pairs, increasing c_d leads to increased overall costs per EV. However, this trend is not true when Myopic is paired with Random or CTMDP2. A possible explanation for this anomaly is that increasing c_d will restrict potential candidate docks under the Myopic policy. We believe Myopic can give guidance when delay costs are high by removing the busier docks from consideration.

An interesting observation is that for $c_d = 200$, Myopic-CTMDP2 out-performs Myopic-Earliest. A scheduling policy using only stochastic information out-performs the policy that exploits exact information with Myopic admission. In the previous experiments, we have seen a large dominance when using observable information over stochastic but clearly this is not always the case. We return to this observation below.

The last experiment studies a charging facility with ten docks and ten cables each. Costs are as in the first two experiment: $c_r = 1$ and $c_d = 5$. We vary the arrival rate $\lambda = \{50, 55, 60\}$ to observe how the policies behave under varying system loads. Figure 4 graphs the results. We see a larger increase in costs for the Free Cable based policy pairs as load increases when contrasted with CTMDP1 and Myopic. Of interest in particular is the performance of Free Cable in comparison to CTMDP1 when using the Earliest scheduling policy. As before, we see at $\lambda = 50$ that Free Cable is better. However, as λ increases, CTMDP1 becomes better.

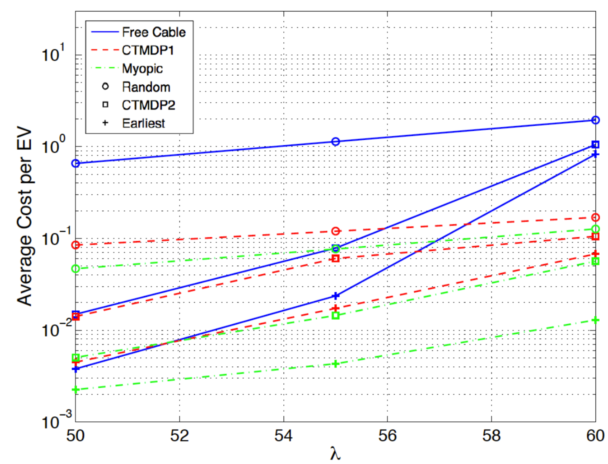

Figure 4: Experiment 4 - Multiple Dock Charging Facility.

6 Discussion and Future Work

The results from simulating the charging facility provide insights into the facility designs as well as the directions for building stronger system management models. For facility design, we see that multi-cable docks provide large performance improvements when there is a disparity between the charging requirements of an EV and the expected deadlines ($p_j < L$). Botsford (2012) discusses such systems and provides two solutions: valet parking and reduced power charging docks to increase charging time. Although these solutions increase dock utilization, they are not always practical as the cost for valet parking or requiring a dock for every customer can be high. If L is much larger than p_j, adding cables to a dock will greatly improve utilization.

Comparisons of the different policy combinations gives us insight into how one would manage a charging facility.

The most important questions that must first be addressed is what information is available and how much control exists for the admission and scheduling of a customer. Parking facilities are used in a variety of settings and different parking lots will have different features. As mentioned in Section 2, management may not have substantial control over a shopping mall parking lot. Customers arrive and choose their spots freely as long as there is space. We can see the similarities to Free Cable-Random where there is no control over the customers. Here, the only decision making required is one of capacity planning; how many docks should be purchased and how many cables will these docks have. As we can see from Figures 1 and 2, more cables may lead to a *decrease* in performance, so choosing the right capacity is very important to the overall system costs. If the manager does have some control over the assignment of EVs and stochastic information is available without observable charging times, CTMDP1-CTMDP2 is the best performing policy. Since exact charging times are not known, using the Earliest scheduling rule is not possible and our experiments show, as expected, that Free Cable does not outperform CTMDP1 in these circumstances.

Although the policies using observable information were able to achieve the best overall performance, our results suggest the potential for using stochastic information. With increased delay costs, the best performing policy combination was Myopic-CTMDP2. This combination makes use of cardinality, stochastic, and observable information. We believe that to achieve the best performance, policies designed to use the stochastic and observable information is required. Myopic-CTMDP2 only achieves the lowest cost in one scenario, but a more sophisticated scheduler that actively uses all the available information has potential to perform favourably on most, if not all, cases.

Such hybrid reasoning in optimization has previously been proposed in the literature. Recent work by Terekhov et al. (2012) and Tran et al. (2013) looks at combining queueing theory and scheduling models to incorporate stochastic reasoning into combinatorial optimization. In Tran et al.'s (2013) work, a queueing model, using stochastic information, was shown to out-perform a number of scheduling models that made use of observable information. The seeming inconsistency with our result is interesting but is likely due to the very different underlying systems and solution approaches. However, the investigation of when combining stochastic and observable information benefits performance is a promising area of future work and our results from testing a system with high delay costs is an example of how such a combination can be advantageous.

Another direction for future work is online stochastic combinatorial optimization (OSCO) (Van Hentenryck and Bent 2006). OSCO creates schedules by generating and optimizing over samples of future arrivals derived from stochastic information. We see it as a promising direction, especially for more complicated scheduling problems.

We would like to expand this work by further examining the CTMDP model and building more sophisticated scheduling models that make use of stochastic reasoning. We believe that a deeper understanding of the characteristics of the optimal CTMDP policy will help provide necessary components one can utilize when creating a sophisticated scheduler that considers both the dynamics of the system and the combinatorial complexities. As well, we would like to explore possible methodologies of solving the CTMDP. Factored representations of an MDP, which uses dynamic Bayesian networks to represent the stochastic decisions of an MDP, are an interesting possibility for being able to solve the CTMDP (Boutilier, Dearden, and Goldszmidt 2000).

7 Related Work on EV Charging

While there are a few studies on EV charging, we are not aware of any work that investigates our parking lot charging scenario with multi-cable docks and dynamically arriving customers. Furthermore, to the best of our knowledge, our paper is the first to study the performance of a multi-cable dock design where the cables in a charging dock are modelled as a limited resource. In all other works, it is either assumed that docks always switch to other cars immediately or the system has unlimited single-cable docks.

Raghavan and Khaligh (2012) examine the effects of EV charging in a smart grid environment. They emphasize the differences between charging methods and time-of-day (evening or night). Li et al. (2011) use dynamic programming to minimize charging costs when electricity prices vary over time. Sioshansi (2012) develops two mixed integer programming models to minimize costs. These works consider the larger power grid problem where people are charging at home rather than in a shared charging facility.

Lee et al. (2011) focus on the problem of a charging station system where there are multiple charging docks and vehicles have different charge lengths, arrival times, and deadlines. In their work, they assume that complete information is known a priori, including the number of EVs. Each vehicle has a different power consumption profile and the objective is to reduce peak power usage over all time periods.

Work on waiting time performance of charging EVs is due to Qin and Zhang (2011). A network of roads is created where nodes represent rest stops to recharge EVs. Drivers are assumed to stop at nodes to charge when required and immediately leave once they are charged. A performance bound is derived and a distributed scheme is proposed which is shown empirically to perform closely to the bound results.

8 Conclusion

We studied scheduling electric vehicles in a charging facility where customers arrive dynamically and tend to park longer than their charge time. Our study considered three information conditions: cardinality, stochastic, and observable. We formulated a CTMDP to optimize the system performance under the first two conditions and demonstrate that it does not scale to realistic-size problems with multiple docks. However, a single-dock version of the CTMDP is tractable. We proposed and numerically evaluated a number of admission and scheduling schemes building on both the single-dock CTMDP and approaches from the scheduling literature under each of the three information conditions. We found that the information available significantly alters the overall performance of the system by limiting the admission and scheduling policies that can be implemented. Thus, it is crucial for any system manager to properly understand the information limitations of his/her system and choose the appropriate methodology to optimize performance.

References

Botsford, C. 2012. The economics of non-residential level 2 EVSE charging infrastructure. In *Proceedings of the 2012 EVS26 International Battery, Hybrid and Fuel Cell Electric Vehicle Symposium*, 1–10.

Boutilier, C.; Dearden, R.; and Goldszmidt, M. 2000. Stochastic dynamic programming with factored representations. *Artificial Intelligence* 121(1):49–107.

Gross, D., and Harris, C. 1998. *Fundamentals of Queueing Theory*. John Wiley & Sons, Inc.

Howard, R. 1960. *Dynamic Programming and Markov Processes*. MIT Press.

Lee, J.; Park, G.; Kim, H.; and Jeon, H. 2011. Fast scheduling policy for electric vehicle charging stations in smart transportation. In *Proceedings of the 2011 ACM Symposium on Research in Applied Computation*, 110–112. ACM.

Li, Z.; Khaligh, A.; and Sabbaghi, N. 2011. Minimum charging-cost tracking based optimization algorithm with dynamic programming technique for plug-in hybrid electric vehicles. In *Proceedings of the Vehicle Power and Propulsion Conference (VPPC), 2011 IEEE*, 1 –6.

Lippman, S. 1975. Applying a new device in the optimization of exponential queueing systems. *Operations Research* 23:687–710.

Pinedo, M. 2008. *Scheduling: theory, algorithms, and systems*. Springer Verlag.

Puterman, M. L. 1994. *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. John Wiley & Sons, Inc.

Qin, H., and Zhang, W. 2011. Charging scheduling with minimal waiting in a network of electric vehicles and charging stations. In *Proceedings of the Eighth ACM international workshop on Vehicular inter-networking*, 51–60. ACM.

Raghavan, S. S., and Khaligh, A. 2012. Impact of plug-in hybrid electric vehicle charging on a distribution network in a smart grid environment. *Innovative Smart Grid Technologies, IEEE PES* 1–7.

Sioshansi, R. 2012. Modeling the impacts of electricity tariffs on plug-in hybrid electric vehicle charging, costs, and emissions. *Operations Research* 60(2):1–11.

Tate, E.; Harpster, M.; and Savagian, P. 2008. *The Electrification of the Automobile: From Conventional Hybrid, to Plug-in Hybrids, to Extended-range Electric Vehicles*. SAE International.

Terekhov, D.; Tran, T. T.; Down, D. G.; and Beck, J. C. 2012. Long-run stability in dynamic scheduling. In *Proceedings of the 22nd International Conference on Automated Planning and Scheduling (ICAPS 2012)*, 261–269.

Tran, T. T.; Terekhov, D.; Down, D. G.; and Beck, J. C. 2013. Hybrid queueing theory and scheduling models for dynamic environments with sequence-dependent setup times. In *Proceedings of the 23nd International Conference on Automated Planning and Scheduling (ICAPS 2013)*, To Appear.

Van Hentenryck, P., and Bent, R. 2006. *Online Stochastic Combinatorial Optimization*. MIT Press.