Finding Minimum-Cost Paths for Electric Vehicles

Timothy M. Sweda and Diego Klabjan
Department of Industrial Engineering and Management Science
Northwestern University
Evanston, Illinois 60208, USA
E-mail: tsweda@u.northwestern.edu, d-klabjan@northwestern.edu

***Abstract*—Modern route-guidance software for conventional gasoline-powered vehicles does not consider refueling since gasoline stations are ubiquitous and convenient in terms of both accessibility and use. The same technology is insufficient for electric vehicles (EVs), however, as charging stations are much more scarce and a suggested route may be infeasible given an EV's initial charge level. Recharging decisions may also have significant impacts on the total travel time and longevity of the battery, which can be costly to replace, so they must be considered when planning EV routes. In this paper, the problem of finding a minimum-cost path for an EV when the vehicle must recharge along the way is modeled as a dynamic program. It is proven that the optimal control and state space are discrete under mild assumptions, and two different solution methods are presented.**

I. Introduction

In the United States, ownership of battery electric vehicles (EVs) is expected to increase considerably over the next several years as the major automakers bring new EV models to market. The limited driving range of EVs renders them most suitable for short-distance commutes, but access to public charging infrastructure would alleviate range anxiety and enable longer trips. However, at least in the short term, charging stations will be relatively scarce, making long-distance commutes difficult to plan.

Aside from the difficulty of planning routes that include visits to charging stations, the charging decisions made at each station introduce further complications. Unlike gasoline refueling, which occurs at the same rate regardless of how full the vehicle's tank is, battery recharging occurs at different rates depending on the initial state of charge. The charging rate is typically highest when the battery's state of charge is low, and it decreases gradually until the battery is fully charged. The longevity of an EV's battery also depends on the charging pattern, including the types of charging stations visited as well as the extent to which the battery is charged. Fast-charging stations that operate at higher voltages can greatly diminish battery life, as can perpetually maintaining a high state of charge in the battery.

In this paper, the problem of finding a minimum-cost path for an EV when the vehicle must recharge along the way is modeled as a dynamic program. It is proven that the optimal control and state space are discrete under mild assumptions, and two different solution methods are presented. One is a backward recursion approach that finds an optimal path when the state space is discrete. The other is an approximate dynamic programming (ADP) algorithm that is applicable to general problem instances.

II. Literature review

The refueling problem for gasoline-powered vehicles, where drivers must decide at which nodes to refuel as well as how much to refuel in order to minimize total fuel costs, has been well studied. It is shown in [1] and [2] that the optimal refueling policy along a fixed path can be solved easily with dynamic programming when fuel prices at each node are static and deterministic. For such a problem, the optimal decision at each node is always one of the following: do not refuel, refuel completely, or refuel just enough to get to the next node where refueling occurs. An algorithm for simultaneously finding the optimal path and refueling policy in a network is detailed in [3], and some combinatorial properties of the optimal policies are explored in [4]. Specifically, it is proven that the problem of finding all-pairs optimal refueling policies reduces to an all-pairs shortest path problem that can be solved in polynomial time.

The objective of the vehicle refueling problem is to minimize the total cost, or the total amount paid for fuel. In reality, however, drivers often want to minimize their travel time as well. While they may be willing to deviate from the shortest path to their destination in order to reduce their fuel expenditures, they will not take major detours for only marginal fuel cost savings. A generic model for vehicle refueling is presented in [5] that attempts to capture such behavior, penalizing longer routes and routes with more refueling stops. The results show that when costs associated with time are included in the objective, significant reductions are achieved in the overall trip cost, fuel used, and miles traveled. Approaches for finding optimal refueling policies when fuel prices are stochastic are given in [6] and [7]. In [6], a forecasting model for predicting future fuel prices is used to generate parameters for a deterministic mixed integer program, and in [7], a dynamic programming framework is presented that is designed to grant drivers greater autonomy to select the stations where they will refuel. These models are difficult to solve analytically, though, and the authors develop only simple heuristics that can be implemented to obtain reasonable solutions.

One assumption that all of the aforementioned models share is that refueling costs are linear. In other words, the cost per unit of fuel at a station is the same regardless of how many units are purchased and how full the vehicle's tank was when it arrived at the station. While this assumption may be true for vehicles powered by liquid fuels such as gasoline, it does not necessarily hold for EVs. Battery charging times are nonlinear, and when the time it takes to recharge is on the order of hours rather than minutes, the value of time during recharging is significant.

A recent thread of research has tackled the problem of finding the most energy-efficient path in a network for EVs. In [8], the authors model the problem as a shortest path problem with constraints on the charge level of the vehicle, such that the charge level can never be negative and cannot exceed the maximum charge level of the battery. Edge weights are permitted to be negative to represent energy recapturing from regenerative braking, but no negative cycles exist. A simple algorithm for solving the problem is provided, and more efficient algorithms are presented in [9] and [10]. It is shown in [9] that the battery capacity constraints can be modeled as cost functions on the edges, and a transformation of the edge cost functions permits the application of Dijkstra's algorithm. The approach described in [10] avoids the use of preprocessing techniques so that edge costs can be calculated dynamically, and it achieves an order of magnitude reduction in the time complexity of the algorithm from [8]. However, none of these methods consider recharging decisions at nodes.

III. Model

Let $G = (N, A)$ be a directed, acyclic network with recharging capability at every node in N. Arc lengths are non-negative and denoted d_{ij} for every arc $(i,j) \in A$, and it is assumed that for each sequence of arcs $\{(i,j),(j,k)\} \subset A$, there exists an arc $(i,k) \in A$ such that $d_{ik} \leq d_{ij} + d_{jk}$ (i.e., the arc lengths satisfy the triangle inequality). The goal is to find the minimum-cost path from s to t $(s,t \in N)$ such that the charge level of the vehicle always remains between a given minimum threshold, h, and the maximum capacity of the battery, q_{max}. The parameter h is permitted to be greater than zero to account for range anxiety and also the fact that the vehicle's performance can be negatively affected when the battery's charge level drops too low.

Let $u_i(q)$ denote the utility of departing from node i with charge level q, where $u_i(\cdot)$ is continuously differentiable, nondecreasing, and defined over the interval $q \in [h, q_{max}]$. If the EV arrives at i with charge level q, then the cost of recharging up to level r is given by the function

$$g_i(q,r) = u_i(r) - u_i(q).$$

Although opportunities may exist for EVs to discharge some of their power back to the grid at charging stations, it is assumed in this paper that $r \geq q$ because an EV driver may not recover the full retail value of electricity discharged from the battery, and discharging also incurs time and battery wear costs. Thus, considering these factors, the actual cost of discharging the battery down to a level $r < q$ would in fact be greater than $u_i(q) - u_i(r)$.

In addition to the previously defined parameters, let f denote the energy efficiency of the EV (in charge per unit distance) and c denote the cost per distance of traversing a link. Both f and c are assumed to be positive constants. Then the optimality equations for the dynamic programming formulation are

$$V(i,q) = \min_{(j,r) \in \mathcal{A}(i,q)} \{g_i(q,r) + cd_{ij} + V(j, r - fd_{ij})\}, \quad (1)$$

where $V(i,q)$ is the value function representing the cost to reach node t from node i with charge level q, and the three terms correspond to the cost of recharging, the cost of traveling from i to j, and the value function upon reaching j, respectively. The feasible action space, $\mathcal{A}(\cdot)$, is defined as

$$\mathcal{A}(i,q) = \left\{ (j,r) : \begin{array}{l} (i,j) \in A, \ h + fd_{ij} \leq q_{max}, \\ r \in [\max\{q, h + fd_{ij}\}, q_{max}] \end{array} \right\}, \quad (2)$$

and the state space, $\mathcal{S}$, is defined as

$$\mathcal{S} = \{(i,q) : i \in N, \ q \in [h, q_{max}]\}. \quad (3)$$

If the vehicle's initial charge level at the origin and desired charge level at the destination are both assumed to be h (i.e., the vehicle should begin and end its trip at the minimum charge level), then the terminal value function is

$$V(t,h) = 0$$

and the objective value is $V(s,h)$.

IV. Algorithms

Because $\mathcal{S}$ is continuous (in the coordinate q), standard recursion methods cannot be implemented directly without first either discretizing the state space or using an approximation algorithm. These approaches do not necessarily yield an optimal solution, but when certain assumptions hold it can be proven that $\mathcal{S}$ reduces to a discrete set, and thus a simple backward recursion algorithm guarantees an optimal solution. This proof and the related algorithm are presented in the next section, followed by an ADP algorithm that can be employed when the assumptions permitting a discrete reduction of $\mathcal{S}$ are relaxed.

A. Backward recursion

In this section, it is proven that under certain assumptions, $\mathcal{S}$ can be discretized without creating an optimality gap. It is first shown that a feasible path can be improved by removing nodes where recharging does not occur, and that there exists an optimal solution where the EV recharges at every visited node (except t). The optimal recharging policy is then proven to be solvable using backward recursion.

Lemma 1: Let $i, n_1, n_2, ..., n_k, j$ be a sequence of nodes visited by a feasible path in G such that the vehicle recharges at i, recharges at j if $j \neq t$, and does not recharge at $n_1, ..., n_k$. Then there exists another feasible path at least as good as the original path that visits j directly from i.

Proof: Since there is no recharging at $n_1, ..., n_k$, the charging cost is zero at those nodes and only the distance cost is considered. It follows from the triangle inequality that

$$d_{ij} \leq d_{i,n_1} + \sum_{\ell=1}^{k-1} d_{n_\ell, n_{\ell+1}} + d_{n_k,j},$$

and thus the total distance cost and recharging cost to reach j from i directly do not exceed the corresponding costs to reach j indirectly. Therefore, there exists a feasible path containing arc $(i,j) \in A$ that is at least as good as the path visiting but not recharging at $n_1, ..., n_k$. ■

If the cost of a feasible path can be decreased by removing nodes where recharging does not occur, then there must exist an optimal path that only visits nodes where recharging does occur. This result is formalized in the following theorem.

Theorem 1: There exists an optimal path from s to t in G such that recharging occurs at every visited node (except t).

Proof: Suppose an optimal path exists that visits but does not recharge at nodes $n_1, ..., n_k$. Then by Lemma 1, there exists another optimal path that bypasses $n_1, ..., n_k$. ■

A consequence of this theorem is that an optimal path can be constructed considering only how much to recharge at each node, not whether or not to recharge. This result alone does not make the problem of finding an optimal recharging policy easier, since the state space is still continuous, but it will simplify the analysis in later parts. In order to reduce the state space to a discrete set, the optimal action space must first be discretized. For arbitrary functions $u_i(\cdot)$, it is not possible to reduce the optimal action space without additional restrictions. However, under mild assumptions, it can be shown that the optimal recharging decision at each node simplifies nicely. Let U_i^+ and U_i^- denote the maximum and minimum gradients, respectively, of $u_i(\cdot)$ for every $i \in N$. They can be expressed mathematically as

$$U_i^+ = \max_{q \in [h, q_{max}]} \{\nabla u_i(q)\}, \quad U_i^- = \min_{q \in [h, q_{max}]} \{\nabla u_i(q)\}.$$

If either $U_i^+ \leq U_j^-$ or $U_j^+ \leq U_i^-$ for any $i, j \in N$, then the optimal action space at each node is a discrete set. The following theorem proves this result.

Theorem 2: Let $m_1, m_2, ..., m_k$ be the sequence of nodes visited by an optimal path where recharging occurs. Then the following is an optimal policy for recharging at each node: At m_k recharge the battery up to level $h + fd_{m_k,t}$; for $\ell = 1, ..., k-1$,

a.) if $U_\ell^- \geq U_{\ell+1}^+$, recharge the battery up to level $h + fd_{m_\ell, m_{\ell+1}}$, and

b.) if $U_\ell^+ \leq U_{\ell+1}^-$, recharge the battery up to level q_{max}.

Proof: Suppose $U_\ell^- \geq U_{\ell+1}^+$ but the optimal solution recharges the battery to a level greater than $h + fd_{m_\ell, m_{\ell+1}}$ at m_ℓ. Then the total cost can be decreased by reducing the amount recharged at m_ℓ and increasing the amount recharged at $m_{\ell+1}$, which contradicts the optimality assumption. Similarly, suppose $U_\ell^+ \leq U_{\ell+1}^-$ but the optimal solution recharges the battery to a level less than q_{max} at m_ℓ. Then the total cost can be decreased by increasing the amount recharged at m_ℓ and reducing the amount recharged at $m_{\ell+1}$, which also contradicts the optimality assumption. Therefore, the stated policy is optimal. ■

From this theorem, it follows that the size of the optimal action space at any node $i \in N$ reduces to at most $2|N-1|$ actions since there are no more than $|N-1|$ other nodes reachable from i and two possible recharging options for each node. Equation (2) can be updated to give the optimal action space,

$$\mathcal{A}^*(i,q) = \left\{ (j,r) : \begin{array}{c} (i,j) \in A, \ h + fd_{ij} \leq q_{max}, \\ r \in \{h + fd_{ij}, q_{max}\} \end{array} \right\}, \tag{4}$$

and modifying (3) yields the new state space,

$$\mathcal{S}^* = \left\{ (i,q) : \begin{array}{l} i \in N, \ q \in \{h\} \cup \\ \qquad \{q_{max} - fd_{ji} \geq h : (j,i) \in A\} \end{array} \right\}, \tag{5}$$

which is a discrete set. Therefore, (1) becomes

$$V(i,q) = \min_{(j,r) \in \mathcal{A}^*(i,q)} \{g_i(q,r) + cd_{ij} + V(j, r - fd_{ij})\} \tag{6}$$

and backward recursion can be implemented to find an optimal path that recharges at every node.

In some special cases, it may be difficult to compute U_i^+ and U_i^- precisely. However, for most practical applications, finding the extreme values of the utility function gradients is manageable, especially when they are monotone. For example, the $u_i(\cdot)$ functions may be convex if they primarily account for time costs, as the charge rate typically decreases at higher charge levels. If $u_i(\cdot)$ happens to be linear and $U_i^+ = U_i^-$ for all $i \in N$, then the problem simplifies to the vehicle refueling problem studied elsewhere in the literature.

B. Approximate dynamic programming

Without the assumptions on the utility functions $u_i(\cdot)$ in the previous section, Theorem 2 no longer holds. In such a scenario, discretization of the state space may not be possible without creating an optimality gap, and thus an approximation algorithm may be appropriate. The following is an ADP algorithm that can be used to solve the vehicle recharging problem.

Step 0: Initialization
- Initialize $\overline{V}^0(i,q)$ for all states $(i,q) \in \mathcal{S}$
- Set initial state: $i^1 = s, \ q^1 = h$
- Set $n = 1$

Step 1: While $(i \neq t)$ loop
- Update value function estimate:

$$\overline{V}^n(i,q) = \begin{cases} \hat{V}^n(i,q), & i = i^n \text{ and } q = q^n \\ \overline{V}^{n-1}(i,q), & \text{otherwise} \end{cases}$$

where

$$\hat{V}^n(i^n,q^n) = \min_{(j^n,r^n) \in \mathcal{A}(i^n,q^n)} \{g_i(q^n,r^n) + cd_{i^n,j^n} + \overline{V}^{n-1}(j^n, r^n - fd_{i^n,j^n})\}$$

- Update current state: $i^n = j^n, \ q^n = r^n - fd_{i^n,j^n}$

Step 2: Let $n = n+1$; if $n \leq iterLimit$, go to Step 1

The main difficulty with this algorithm is the initialization of the value function estimates. The performance of the algorithm is sensitive to these initial values. Further investigation is needed to find good starting estimates for the value functions that ensure quality solutions.

V. Conclusions and Future Work

In this paper, the problem of finding a minimum-cost path for an EV when the vehicle must recharge along the way is modeled as a dynamic program. It is proven that the optimal state space can be discretized when certain conditions are satisfied, and two algorithms for finding an optimal path are proposed. The first is a backward recursion method that can be applied when the state space is discrete. The second is an ADP algorithm that is useful for more general instances of the model.

As the vehicle recharging problem is still fairly new in the literature, there are many directions for future work. One unique aspect of EVs not captured in this paper is the ability to recapture energy through regenerative braking. By traveling downhill, an EV can sometimes increase its battery's charge level while braking, leading to cases where the net energy consumption along a link is negative. Time could also be introduced as another dimension to the model to permit dynamic and stochastic utilities of recharging. Such models would inevitably be more complex, necessitating the development of robust approximation schemes such as ADP.

References

[1] S. H. Lin et al, "A linear-time algorithm for finding optimal vehicle refueling policies," in *Operations Research Lett.*, 2007, vol. 35, no. 3, pp. 290-296.

[2] S. Khuller et al., "To fill or not to fill: the gas station problem," in *ACM Trans. Algorithms*, 2011, vol. 7, no. 3, pp. 1-16.

[3] S. H. Lin, "Finding optimal refueling policies: a dynamic programming approach," in *J. Computing Sci. Colleges*, 2008, vol. 23, no. 6, pp. 272-279.

[4] S. H. Lin, "Finding optimal refueling policies in transportation networks," in *Algorithmic Aspects Inform. Manage.*, 2008, vol. 5034, pp. 280-291.

[5] Y. Suzuki, "A generic model of motor-carrier fuel optimization," in *Naval Research Logistics*, 2008, vol. 55, no. 8, pp. 737-746.

[6] E. Klampfl et al., "Intelligent refueling advisory system," in *Proc. Int. Conf. Informatics Control Automation Robotics*, Madeira, Portugal, 2008.

[7] Y. Suzuki, "A decision support system of dynamic vehicle refueling," in *Decision Support Syst.*, 2009, vol. 46, no. 2, pp. 522-531.

[8] A. Artmeier et al., "The shortest path problem revisited: optimal routing for electric vehicles," in *KI*, 2010, vol. 6359, pp. 309-316.

[9] J. Eisner et al., "Optimal route planning for electric vehicles in large networks," in *Proc. 25th Assoc. Advancement Artificial Intell. Conf.*, San Francisco, CA, 2011.

[10] M. Sachenbacher et al., "Efficient energy-optimal routing for electric vehicles," in *Proc. 25th Assoc. Advancement Artificial Intell. Conf.*, San Francisco, CA, 2011.