On the interaction between Autonomous Mobility-on-Demand systems and the power network: models and coordination algorithms

Federico Rossi, Ramon Iglesias, Mahnoosh Alizadeh, and Marco Pavone

Abstract—This paper studies the interaction between a fleet of electric, self-driving vehicles servicing on-demand transportation requests (referred to as Autonomous Mobility-on-Demand, or AMoD, system) and the electric power network. We propose a *joint* linear model that captures the coupling between the two systems stemming from the vehicles' charging requirements. The model subsumes existing network flow models for AMoD systems and DC models for the power network, and it captures time-varying customer demand and power generation costs, road congestion, and power transmission and distribution constraints. We then leverage the model to jointly optimize the operation of both systems. We devise an algorithmic procedure to losslessly reduce the problem size by bundling customer requests, allowing it to be efficiently solved by off-the-shelf linear programming solvers. We then study the implementation of a hypothetical electric-powered AMoD system in Dallas-Fort Worth, and its impact on the Texas power network. We show that coordination between the AMoD system and the power network can *reduce* the overall energy expenditure compared to the case where no cars are present (despite the increased demand for electricity) and yield savings of $182M/year compared to an uncoordinated scenario. Finally, we provide a closed-loop receding-horizon implementation. Collectively, the results of this paper provide a first-of-a-kind characterization of the interaction between electric-powered AMoD systems and the power network, and shed additional light on the economic and societal value of AMoD.

I. INTRODUCTION

Private vehicles are major contributors to urban pollution, which is estimated to cause over seven million premature deaths worldwide every year [29]. Plug-in electric vehicles (EVs) hold promise to significantly reduce urban pollution, both by reducing carbon dioxide emissions from internal-combustion engine vehicles, and by enabling use of renewable and low-polluting power generators as a source of energy for transportation services. However, at present, adoption of EVs for private mobility has been significantly hampered by customers' concerns about limited range and availability of charging infrastructure.

The emerging technology of self-driving vehicles might provide a solution to these challenges and thus might represent a key enabler for the widespread adoption of EVs. Specifically,

fleets of self-driving vehicles providing on-demand transportation services (referred to as Autonomous Mobility-on-Demand, or AMoD, systems) hold promise to replace personal transportation in large cities by offering high quality of service at lower cost [22] with positive effects on safety, parking infrastructure, and congestion. Crucially, EVs are especially well-suited to AMoD systems. On the one hand, short-range trips typical of urban mobility are well-suited to the current generation of range-limited EVs; on the other hand, intelligent fleet-wide policies for rebalancing and charging can ensure that vehicles with an adequate level of charge are available to customers, virtually eliminating "range anxiety," a major barrier to EV adoption. To fully realize this vision, however, one needs currently unavailable tools to manage the complex *couplings* between AMoD fleet management (e.g., for routing and charging the EVs) and the control of the power network. Specifically, one should consider

1) *Impact of transportation network on power network*: Concurrent charging of large numbers of EVs can have significant effects both on the stability of the power network and on the local price of electricity (including at the charging stations) [21, 3, 9]. For example, [9] shows that in California a 25% market penetration of (non-autonomous) EVs with fast chargers, in the absence of smart charging algorithms, would increase overall electricity demand in peak load by about 30%, and electricity prices by almost 200%.

2) *Impact of power network on transportation network*: Electricity prices can significantly affect travel patterns for EVs. [3] shows that changes in electricity prices can radically alter the travel patterns and charging schedules of fleets of EVs in a simplified model of the San Francisco Bay Area. This, in turn, would affect electricity prices in a complex feedback loop.

The key idea behind this paper is that, by intelligently routing fleets of autonomous EVs and, in particular, by harnessing the flexibility offered by the routes and schedules for the empty-traveling vehicles, one can *actively* control such complex couplings and guarantee high-performance for the overall system (e.g., high passenger throughput and lower electricity costs). Additionally, autonomous EVs provide a unique opportunity for joint traffic and energy production management, as they could act as mobile storage devices. That is, when not used for the fulfillment of trip requests, the vehicles could be routed to target charging stations in order to either absorb excess generated energy at time of low power demand (by charging) or inject power in the power network

at times of high demand (by discharging).

Literature review: The integration of *non-autonomous* EVs within the power network has been addressed in three main lines of work. A first line of work addresses the problem of scheduling charging of EVs (i.e., optimizing the charging profile in *time*) under the assumption that the vehicles' charging schedule has no appreciable effect on the power network [20, 24, 25]. This assumption is also commonly made when selecting the locations of charging stations (i.e., optimizing the charging profile in *space*) [8, 18]. A high penetration of EVs would, however, significantly affect the power network. Thus, a second line of work investigates the effects of widespread adoption of EVs on key aspects such as wholesale prices and reserve margins, for example in macroeconomic [9] and game-theoretical [21, 27] settings. Accordingly, [2, 3] investigate joint models for EV routing and power generation/distribution aimed at driving the system toward a socially-optimal solution. Finally, a third line of work investigates the potential of using EVs to regulate the power network and satisfy short-term spikes in power demand. The macroeconomic impact of such schemes (generally referred to as Vehicle-To-Grid, or V2G) has been studied in [12], where it is shown that widespread adoption of EVs and V2G technologies could foster significantly increased adoption of wind power. Going one step further, [13] proposes a unified model for EV fleets and the power network, and derives a joint dispatching and routing strategy that maximizes social welfare (i.e., it minimizes the *overall* cost borne by all participants, as opposed to maximizing individual payoffs). However, [12] does not capture the *spatial* component of the power and transportation networks, while [13] assumes that the vehicles' schedules are fixed.

The objective of this paper is to investigate the interaction between AMoD and electric power systems (jointly referred to as Power-in-the-loop AMoD, or P-AMoD, systems) in terms of modeling and algorithmic tools to effectively manage their couplings. In this context, our work improves upon the state of the art (in particular, [2, 3]) along three main dimensions: (i) it considers a fleet of *shared* and *autonomous* EVs, which offer significant additional degrees of freedom for vehicle scheduling, routing, and charging; (ii) it provides efficient algorithms that can solve large-scale problems; and (iii) it characterizes the vehicles' ability to return power to the power network through V2G schemes, and its economic benefits.

Statement of contributions: Specifically, the contribution of this paper is fourfold. First, we propose a joint linear model for P-AMoD systems. The model captures time-varying customer demand and generation prices, congestion in the road network, power transmission constraints on the transmission lines, and transformer capacity constraints induced by the distribution network. Second, we leverage the model to design tools that optimize the operations of P-AMoD systems and, in particular, maximize social welfare. To this end, we propose an algorithmic procedure to losslessly reduce the dimensionality of the P-AMoD model. The procedure allows P-AMoD problems with hundreds of road links, time horizons of multiple hours, and any number of customers and vehicles to be solved on commodity hardware. Third, we apply the model and algorithms to a case study of a hypothetical deployment of an AMoD system in Dallas-Fort Worth, TX. We show that

coordination between the AMoD system and the electric power network can have a significant positive impact on the price of electricity (remarkably, the overall electricity expenditure in presence of the AMoD system can be *lower* than in the case where no vehicles are present, despite the increased demand), while retaining *all* the convenience and sustainability benefits of AMoD. This suggests that the societal value of AMoD systems spans beyond mobility: properly coordinated, AMoD systems can deliver significant benefits to the wider community by helping increase the efficiency of the power network. Finally, we present a receding-horizon algorithm for P-AMoD that provides built-in closed-loop robustness and delivers computation times in the order of seconds on commodity hardware at the price of some suboptimality.

Organization: The remainder of this paper is organized as follows. In Section II we present a linear model that captures the interaction between an AMoD system and the power network. In Section III, we propose a procedure to losslessly reduce the size of the model by bundling customer requests. In Section IV, we evaluate our model and algorithm on a case study of Dallas-Fort Worth. In Section V, we propose a receding-horizon implementation. Finally, in Section VI, we draw conclusions and discuss directions for future work.

II. MODEL DESCRIPTION AND PROBLEM FORMULATION

We propose a linear, flow-based model that captures the interaction between an AMoD system and the power network. The model consists of two parts.

First, we extend the model in [30] to a time-varying network flow model of an AMoD system with EVs. We assume that a Transportation Service Operator (TSO) manages the AMoD system in order to fulfill passenger trip requests within a given road network. Road links are subject to congestion, and trip requests arrive according to an *exogenous* dynamical process. The TSO must not only compute the routes for the autonomous EVs (i.e. *vehicle routing*), but also issue tasks and routes for empty vehicles in order to realign the fleet with the asymmetric distribution of trip demand (i.e. *vehicle rebalancing*). Due to limited battery capacity, the EVs need to periodically charge at charging stations. The price of electricity varies between charging stations – the charging schedule is determined by the TSO in order to minimize the fleet's operational cost.

The price of electricity itself is a result of the power system operation to balance supply and demand, and varies across the power grid. Thus, we next review the linear (DC) power flow model of the power network and the economic dispatch problem. The power transmission network comprises energy providers that are connected to load buses through high-voltage transmission lines. Transmission capacities (dictated chiefly by thermal considerations) limit the amount of power that can be transferred on each transmission line. Load buses are connected to charging stations and other sources of power demand through the distribution systems: this system induces constraints on the amount of power that can be served to each load bus. Power demands other than those from charging stations are regarded as exogenous parameters in this paper. The power network is controlled by an Independent System Operator (ISO). The ISO also determines prices at the load buses (and, consequently, at the charging stations) so as to

guarantee grid reliability while minimizing the overall generation cost (a problem known as *economic dispatch*).

The vehicles' charging introduces a coupling between the transportation and the power networks. The loads due to charging influence the local price of electricity set by the ISO – the prices, in turn, affect the optimal charging schedule computed by the TSO. Accordingly, we conclude this section by describing the interaction between the two models, and we propose a joint model for Power-in-the-loop AMoD.

A. Network Flow Model of an AMoD system

We consider a time-varying, finite-horizon model. The time horizon of the problem is discretized in T time intervals, each corresponding to T_S seconds; the battery charge level of the autonomous vehicles is similarly discretized in C charge levels, each corresponding to J_C joules.

Road network: The road network is modeled as a directed graph $R = (\mathcal{V}_R, \mathcal{E}_R)$, where $\mathcal{V}_R$ denotes the node set and $\mathcal{E}_R \subseteq \mathcal{V}_R \times \mathcal{V}_R$ denotes the edge set. Nodes $v \in \mathcal{V}_R$ denote either an intersection, a charging station, or a trip origin/destination. Edges $(v, w) \in \mathcal{E}_R$ denote the availability of a road link connecting nodes v and w. For each edge, the length $d_{(v,w)} \in \mathbb{R}_{\geq 0}$ determines the mileage driven along the road link; the traversal time $t_{(v,w)} \in \{1, \dots, T\}$ characterizes the travel time on the road link in absence of congestion; the energy requirement $c_{(v,w)} \in \{-C, \dots, C\}$ models the energy consumption (i.e., the number of charge levels) required to traverse the link in absence of congestion; and the capacity $\overline{f}_{v,w} \in \mathbb{R}_{\geq 0}$ captures the maximum vehicle flow rate (i.e., the number of vehicles per unit of time) that the road link can accommodate without experiencing congestion.

Vehicles traversing the road network can recharge and discharge their batteries at charging stations, whose locations are modeled as a set of nodes $\mathcal{S} \subset \mathcal{V}_R$. The charging and discharging rates $\delta c_s^+, \delta c_s^- \in \{1, \dots, C\}$ correspond to the amount of energy (in charge levels) that the charger can provide to a vehicle (or, conversely, that a vehicle can return to the power grid) in one unit of time. For simplicity, we assume that the charging rates are fixed; however, the model can be extended to accommodate variable charging rates. The charging and discharging prices $p_s^+(t), p_s^-(t) \in \mathbb{R}$ capture the unit cost of energy (or, conversely, the unit payment the vehicles receive for returning power to the grid) at time t; in this paper, $p_s^+(t) = p_s^-(t)$. The capacity $\overline{S}_s \in \mathbb{N}$ models the maximum number of vehicles that can simultaneously charge or discharge at station s.

Expanded AMoD network: We introduce an *expanded AMoD network* modeled as a directed graph $G = (\mathcal{V}, \mathcal{E})$ to represent the time-varying nature of the problem and the state-of-charge of the vehicles. Specifically, each node $\mathbf{v} \in \mathcal{V}$ models a physical location at a given time and charge level, while edges $e \in \mathcal{E}$ model road links and charging actions at a given time and charge level. Formally, a node $\mathbf{v} \in \mathcal{V}$ corresponds to a tuple $\mathbf{v} = (v_{\mathbf{v}}, t_{\mathbf{v}}, c_{\mathbf{v}})$, where $v_{\mathbf{v}} \in \mathcal{V}_R$ is a node in the road network graph R; $t_{\mathbf{v}} \in \{1, \dots, T\}$ is a discrete time; and $c_{\mathbf{v}} \in \{1, \dots, C\}$ is a discrete charge level. The edge set $\mathcal{E}$ is partitioned into two disjoint subsets, namely $\mathcal{E}_L$ and $\mathcal{E}_S$. Edges $e \in \mathcal{E}_L$ represent road links, whereas edges $e \in \mathcal{E}_S$ model the charging/discharging process at the stations. An edge $(\mathbf{v}, \mathbf{w})$ belongs to $\mathcal{E}_L$ when (i) an edge $(v_{\mathbf{v}}, v_{\mathbf{w}})$ exists in the road network graph edge set $\mathcal{E}_R$, (ii) the link $(v_{\mathbf{v}}, v_{\mathbf{w}}) \in \mathcal{E}_R$ can be traversed in time $t_{\mathbf{w}} - t_{\mathbf{v}} = t_{(v_{\mathbf{v}}, v_{\mathbf{w}})}$, and (iii) the battery charge required to traverse the link is $c_{\mathbf{v}} - c_{\mathbf{w}} = c_{(v_{\mathbf{v}}, v_{\mathbf{w}})}$. Conversely, an edge $(\mathbf{v}, \mathbf{w})$ represents a charging/discharging edge in $\mathcal{E}_S$ when (i) $v_{\mathbf{v}} = v_{\mathbf{w}}$ is the location of a charging station in $\mathcal{S}$ and (ii) the charging/discharging rate at the charging location $v_{\mathbf{v}}$ is $(c_{\mathbf{w}} - c_{\mathbf{v}})/(t_{\mathbf{w}} - t_{\mathbf{v}}) = \delta c_{v_{\mathbf{v}}}^+$ (charging) or $(c_{\mathbf{w}} - c_{\mathbf{v}})/(t_{\mathbf{w}} - t_{\mathbf{v}}) = \delta c_{v_{\mathbf{v}}}^-$ (discharging). Figure 1 (left) shows a graphical depiction of the graph G.
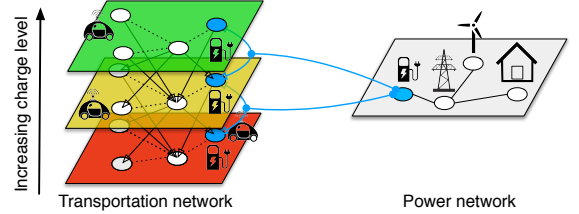

Fig. 1. Augmented transportation and power networks. As vehicles travel on road links (modeled by solid black arrows), their charge level decreases. Blue nodes represent charging stations: the flows on charging and discharging edges affect the load at the corresponding nodes in the power network. For simplicity, only one time step is shown.

Customer and rebalancing routes: Transportation requests are represented by the set of tuples $\{(v_m, w_m, t_m, \lambda_m)\}_{m=1}^M$, where $v_m \in \mathcal{V}_R$ is the request's origin location, $w_m \in \mathcal{V}_R$ is the request's destination location, t_m is the requested pickup time, and λ_m is the average customer arrival rate (or simply customer rate) of request m at time interval t_m. Transportation requests are assumed to be known and deterministic.

The goal of the TSO is to compute a routing and recharging policy for the self-driving vehicles. To achieve this, we model vehicle routes as network flows [1]. Network flows are an *equivalent representation* for routes. Indeed, any route can be represented as a network flow assuming value 1 on edges belonging to the route and 0 elsewhere; conversely, all network flows considered in this paper can be represented as a collection of weighed routes [1, Ch. 3]. This representation allows us to leverage the rich theory of network flows: in particular, in Section III section we exploit this theory to *losslessly reduce* the dimensionality of the optimization problem.

We denote the *customer flow* as the rate of customer-carrying vehicles belonging to a specific transportation request $(v_m, w_m, t_m, \lambda_m)$ traversing an edge $e \in \mathcal{E}$. Formally, for request $m \in \{1, \dots, M\}$, the customer flow is a function $f_m(\mathbf{v}, \mathbf{w}) : \mathcal{E} \mapsto \mathbb{R}_{\geq 0}$, that represents the rate of customers belonging to request m traveling from location $v_{\mathbf{v}}$ to location $v_{\mathbf{w}}$ (or charging/discharging at location $v_{\mathbf{v}} = v_{\mathbf{w}}$) from time $t_{\mathbf{v}}$ to time $t_{\mathbf{w}}$, with an initial battery charge of $c_{\mathbf{v}}$ and a final battery charge of $c_{\mathbf{w}}$. Analogously, the rebalancing (or customer-empty) flow $f_0(\mathbf{v}, \mathbf{w}) : \mathcal{E} \mapsto \mathbb{R}_{\geq 0}$ represents the rate of empty vehicles traversing a road link or charging/discharging. Customer flows must satisfy a *continuity* condition: customer-carrying vehicles entering a node at a given time and charge level must exit the same node at the same time and with the same charge level. Equation (1) enforces this condition:

$$\sum_{\mathbf{u}:(\mathbf{u},\mathbf{v})\in\mathcal{E}} f_m(\mathbf{u}, \mathbf{v}) + 1_{v_{\mathbf{v}}=v_m} 1_{t_{\mathbf{v}}=t_m} \lambda_m^{c_{\mathbf{v}},\text{in}} = \sum_{\mathbf{w}:(\mathbf{v},\mathbf{w})\in\mathcal{E}} f_m(\mathbf{v}, \mathbf{w}) +$$
$$1_{v_{\mathbf{v}}=w_m} \lambda_m^{t_{\mathbf{v}},c_{\mathbf{v}},\text{out}} \quad \forall \mathbf{v} \in \mathcal{V}, m \in \{1, \dots, M\}, \tag{1a}$$

$$\sum_{c=1}^C \lambda_m^{c,\text{in}} = \lambda_m, \quad \sum_{t=1}^T \sum_{c=1}^C \lambda_m^{t,c,\text{out}} = \lambda_m, \quad \forall m \in \{1, \dots, M\}, \tag{1b}$$

where the variable $\lambda_m^{c,\text{in}}$ denotes the customer rate departing with charge level c and the variable $\lambda_m^{t,c,\text{out}}$ denotes the customer rate reaching the destination at time t with charge level c; both are optimization variables. Function 1_x denotes the indicator function of the Boolean variable $x = \{\text{true}, \text{false}\}$, that is $1_x = 1$ if x is true, and $1_x = 0$ if x is false.

Rebalancing flows must satisfy a continuity condition analogous to the one for the customer flows. In addition, rebalancing flows must satisfy a *consistency* condition representing the fact that a customer may only depart the origin location if an empty vehicle is available. Finally, the initial position and charge level of the vehicles is fixed; the final position and charge level is an optimization variable (possibly subject to constraints, e.g., on the minimum final charge level). The constraints for the initial and final positions of the rebalancing vehicles at each node $\mathbf{v} \in \mathcal{V}$ are captured by a set of functions $N_I(\mathbf{v})$ and $N_F(\mathbf{v})$, respectively. Formally, $N_I(\mathbf{v})$, with $t_{\mathbf{v}} = 1$, denotes the number of rebalancing vehicles entering the AMoD system at location $v_{\mathbf{v}}$ with charge level $c_{\mathbf{v}}$. Conversely, $N_F(\mathbf{v})$, with $t_{\mathbf{v}} = T$ denotes the number of rebalancing vehicles at location $v_{\mathbf{v}}$ with charge level $c_{\mathbf{v}}$. For $t_{\mathbf{v}} \neq 1$, $N_I(\mathbf{v}) = 0$; for $t_{\mathbf{v}} \neq T$, $N_F(\mathbf{v}) = 0$. The overall number of vehicles in the network is $\sum_{\mathbf{v} \in \mathcal{V}} N_I(\mathbf{v})$. Equation (2) simultaneously enforces the rebalancing vehicles' continuity condition, consistency condition, and the constraints on the initial and final locations:

$$\sum_{\mathbf{u}:(\mathbf{u},\mathbf{v})\in\mathcal{E}} f_0(\mathbf{u},\mathbf{v}) + \sum_{m=1}^{M} 1_{v_{\mathbf{v}}=w_m} \lambda_m^{t_{\mathbf{v}},c_{\mathbf{v}},\text{out}} + N_I(\mathbf{v}) =$$

$$\sum_{\mathbf{w}:(\mathbf{v},\mathbf{w})\in\mathcal{E}} f_0(\mathbf{v},\mathbf{w}) + \sum_{m=1}^{M} 1_{v_{\mathbf{v}}=v_m} 1_{t_{\mathbf{v}}=t_m} \lambda_m^{c_{\mathbf{v}},\text{in}} + N_F(\mathbf{v}), \forall \mathbf{v} \in \mathcal{V}. \quad (2)$$

Congestion: We adopt a simple *threshold* model for congestion: the vehicle flow on each road link is constrained to be smaller than the road link's capacity. The model is analogous to the one adopted in [30] and is consistent with classical traffic flow theory [28]. This simplified congestion model is adequate for our goal of *controlling* the vehicles' routes and charging schedules, and ensures tractability of the resulting optimization problem; higher-fidelity models can be used for the *analysis* of the AMoD system's operations. Equation (3) enforces the road congestion constraint:

$$\sum_{c_{\mathbf{v}}=1}^{C} \sum_{m=0}^{M} f_m(\mathbf{v},\mathbf{w}) \leq \overline{f}_{(v_{\mathbf{v}},v_{\mathbf{w}})}, \forall (v_{\mathbf{v}},v_{\mathbf{w}}) \in \mathcal{E}_R, t_{\mathbf{v}} \in \{1,\ldots,T\}$$

$$(3)$$

Charging stations can simultaneously accommodate a limited number of vehicles. The station capacity constraint is enforced with Equation (4):

$$\sum_{\substack{(\mathbf{v},\mathbf{w})\in\mathcal{E}_S:\\v_{\mathbf{v}}=v_{\mathbf{w}}=v}} \sum_{m=0}^{M} f_m(\mathbf{v},\mathbf{w}) \leq \overline{S}_{v_{\mathbf{v}}}, \forall v \in \mathcal{S}, t \in \{1,\ldots,T\}. \quad (4)$$

Network flow model of an AMoD system: The travel time T_M experienced by customers, a proxy for customer welfare, and the overall mileage D_V driven by (both customer-carrying and empty) vehicles, a proxy for vehicle wear, are given by

$$T_M = \sum_{(\mathbf{v},\mathbf{w})\in\mathcal{E}} t_{\mathbf{v},\mathbf{w}} \sum_{m=1}^{M} f_m(\mathbf{v},\mathbf{w}),$$

$$D_V = \sum_{(\mathbf{v},\mathbf{w})\in\mathcal{E}} d_{v_{\mathbf{v}},v_{\mathbf{w}}} \sum_{m=0}^{M} f_m(\mathbf{v},\mathbf{w}),$$

Note that T_M only includes the travel time of *customer-carrying* vehicles, whereas D_V includes the distance traveled by *all* vehicles. Also note that, for charging edges, $d_{v_{\mathbf{v}},v_{\mathbf{w}}} = 0$. The overall cost of electricity incurred by the vehicles (including credits from selling electricity to the power network) is

$$V_E = \sum_{(\mathbf{v},\mathbf{w})\in\mathcal{E}_S} \sum_{m=0}^{M} f_m(\mathbf{v},\mathbf{w}) \delta c_{v_{\mathbf{v}}} p_{(\mathbf{v},\mathbf{w})},$$

where $\delta c_{v_{\mathbf{v}}} = \delta c_{v_{\mathbf{v}}}^+$ and $p_{(\mathbf{v},\mathbf{w})} = p_{v_{\mathbf{v}}}^+$ if $c_{\mathbf{w}} > c_{\mathbf{v}}$, $\delta c_{v_{\mathbf{v}}} = \delta c_{v_{\mathbf{v}}}^-$ and $p_{(\mathbf{v},\mathbf{w})} = p_{v_{\mathbf{v}}}^-$ otherwise.

The goal of the TSO is to solve the Vehicle Routing and Charging problem, that is, to minimize the aggregate societal cost borne by the AMoD users while satisfying all operational constraints. We define the customers' value of time (i.e., the monetary loss associated with traveling for one unit of time) as V_T and the operation cost per kilometer of the vehicles (including maintenance but excluding electricity costs) as V_D. The aggregate societal cost experienced by the AMoD users is then defined as

$$V_D D_V + V_E + V_T T_M. \quad (5)$$

The Vehicle Routing and Charging problem entails minimizing (5) subject to constraints (1), (2), (3), and (4).

B. Linear model of power network

In this paper, the power network is modeled according to the well-known DC model [14, Ch. 6], which, by assuming constant voltage magnitudes and determining the power flow on transmission lines solely based on voltage phase angles, represents an approximation to the higher-fidelity AC flow model [7]. In analogy with the treatment of the AMoD model, we discretize the time horizon of the problem in T time steps. The power grid is modeled as an undirected graph $P = (\mathcal{B}, \mathcal{E}_P)$, where $\mathcal{B}$ is the node set, commonly referred to as buses in the power engineering literature, and $\mathcal{E}_P \subseteq \mathcal{B} \times \mathcal{B}$ is the edge set, representing the transmission lines. The subsets of buses representing generators and loads are defined as $\mathcal{G} \subset \mathcal{B}$ and $\mathcal{L} \subset \mathcal{B}$, respectively. Nodes that are neither loads nor generators are referred to as interconnects. Generators produce power and deliver it to the network, while loads absorb power from the network. Each generator $g \in \mathcal{G}$ is characterized by a maximum output power $\overline{p}_g(t)$, a minimum output power $\underline{p}_g(t)$, a unit generation cost $o_g(t)$, and maximum ramp-up and ramp-down rates $p_g^+(t)$ and $p_g^-(t)$, respectively. Transmission lines $e \in \mathcal{E}_P$ are characterized by a reactance x_e and a maximum allowable power flow $\overline{p}_e$ (due chiefly to thermal constraints). The reactance and the maximum allowable power flow do not vary with time. Each load node $l \in \mathcal{L}$ is characterized by a required power demand $d_l(t)$. The distribution network is not modeled explicitly; however, thermal constraints due to the distribution substation transformers are modeled by an upper bound $\overline{d}_l(t)$ on the power that can be delivered at each load node.

We define a generator power function $p : (\mathcal{G}, \{1,\ldots,T\}) \mapsto \mathbb{R}_{\geq 0}$, and a phase angle function $\theta : (\mathcal{B}, \{1,\ldots,T\}) \mapsto \mathbb{R}$. The Economic Dispatch problem entails minimizing the generation cost subject to a set of feasibility constraints [14], namely:

$$\underset{p,\theta}{\text{minimize}} \quad \sum_{t=1}^{T}\sum_{g\in\mathcal{G}} o_g(t)p(g,t), \tag{6a}$$

$$\text{subject to} \sum_{(u,v)\in\mathcal{E}_P} \frac{\theta(u,t)-\theta(v,t)}{x_{u,v}} + 1_{v\in\mathcal{G}}p(v,t) = 1_{v\in\mathcal{L}}d_v(t)+$$

$$\sum_{(v,w)\in\mathcal{E}_P} \frac{\theta(v,t)-\theta(w,t)}{x_{v,w}}, \forall v\in\mathcal{B}, t\in\{1,\ldots,T\}, \tag{6b}$$

$$-\overline{p}_{b_1,b_2} \le \frac{\theta(b_1,t)-\theta(b_2,t)}{x_{b_1,b_2}} \le \overline{p}_{b_1,b_2},$$
$$\forall(b_1,b_2)\in\mathcal{E}_P, t\in\{1,\ldots,T\}, \tag{6c}$$

$$\underline{p}_g(t) \le p(g,t) \le \overline{p}_g(t), \forall g\in\mathcal{G}, t\in\{1,\ldots,T\}, \tag{6d}$$

$$-p_g^-(t) \le p(g,t+1)-p(g,t) \le p_g^+(t),$$
$$\forall g\in\mathcal{G}, t\in\{1,\ldots,T-1\}, \tag{6e}$$

$$d_l(t) \le \overline{d}_l(t), \qquad \forall l\in\mathcal{L}, t\in\{1,\ldots,T\}. \tag{6f}$$

Equation (6b) enforces power balance at each bus based on the DC power flow equations; Equation (6c) encodes the transmission lines' thermal constraints; Equation (6d) captures the generation capacity constraints; Equation (6e) encodes the ramp-up and ramp-down constraints; and Equation (6f) encodes the thermal constraints of substation transformers at load nodes.

The unit price of electricity at the load nodes is typically determined through a mechanism known as Location Marginal Pricing (LMP) [14]. The LMP at a node is defined as the *marginal cost* of delivering one unit of power at the node while respecting all the system constraints. It can be shown [14] that the LMP at each bus equals the dual variable corresponding to the power injection constraint (6b) at the same bus in the Economic Dispatch problem.

C. Power-in-the-loop AMoD system

The vehicles' charging requirements introduce a *coupling* between the AMoD system and the power network, as shown in Figure 1. Specifically, the vehicles' charging schedule produces a load on the power network. Such a load on the power network affects the solution to the ISO's Economic Dispatch problem and, as a result, the LMPs. The change in LMPs, in turn, has an effect on the TSO's optimal charging schedule. In absence of coordination, this feedback loop can lead to system instability, as shown for the case of privately-owned, non-autonomous EVs in [3].

In this section, we formulate a *joint model* for the TSO's Vehicle Routing and Charging problem and the ISO's Economic Dispatch problem. The goal of this model is to maximize the social welfare by minimizing the total cost of mobility (a profit-maximizing formulation would be similar) and the total cost of power generation and transmission. While the resulting solution is not directly actionable (since it requires the TSO and the ISO to coordinate and share their private information), coordination mechanisms can be designed to steer the system towards the optimum: specifically, in the extended version of this paper [19] we show that the solution can be enforced as a general equilibrium through appropriate pricing.

The coupling between the AMoD model and the electric power model is mediated by the charging stations. A given charging station is represented both by a node $v \in \mathcal{V}_R$

in the road network and by a load node $l \in \mathcal{L}$ in the power network. To capture this correspondence, we define an auxiliary function $\mathcal{M}_{\text{P,R}} : \mathcal{L} \mapsto \{\mathcal{V}_R \cup \emptyset\}$. Given a load node $b \in \mathcal{L}$, $\mathcal{M}_{P,R}(b)$ denotes the node in $\mathcal{V}_R$ (if any) that represents a charging station connected to b. We then define two additional functions, $\mathcal{M}_{\text{P,G}}^+ : (\mathcal{L}, \{1,\ldots,T\}) \mapsto \{\mathcal{E} \cup \emptyset\}$ and $\mathcal{M}_{\text{P,G}}^- : (\mathcal{L}, \{1,\ldots,T\}) \mapsto \{\mathcal{E} \cup \emptyset\}$. The function $\mathcal{M}_{\text{P,G}}^+$ (resp. $\mathcal{M}_{\text{P,G}}^-$) maps a load node l and a time t to the set of charge (resp. discharge) edges in G corresponding to station $\mathcal{M}_{\text{P,R}}(l)$ at time t. Formally,

$$\mathcal{M}_{\text{P,G}}^+(l,t) :\{(\mathbf{v},\mathbf{w})\in\mathcal{E}_S | v_{\mathbf{v}}\in\mathcal{M}_{\text{P,R}}(l), c_{\mathbf{v}} < c_{\mathbf{w}}, t_{\mathbf{v}} \le t < t_{\mathbf{w}}\},$$
$$\mathcal{M}_{\text{P,G}}^-(l,t) :\{(\mathbf{v},\mathbf{w})\in\mathcal{E}_S | v_{\mathbf{v}}\in\mathcal{M}_{\text{P,R}}(l), c_{\mathbf{v}} > c_{\mathbf{w}}, t_{\mathbf{v}} \le t < t_{\mathbf{w}}\}.$$

The load at a load bus l can be expressed as the sum of two components: an exogenous demand $d_{l,e}$ and the load due to the charger or chargers connected to that bus, quantitatively,

$$d_l(t) = d_{l,e}(t) + J_C\delta c_{\mathcal{M}_{\text{P,R}}(l)}^+ \sum_{(\mathbf{v},\mathbf{w})\in M_{P,G}^+(l,t)} \sum_{m=0}^{M} f_m(\mathbf{v},\mathbf{w})$$
$$+ J_C\delta c_{\mathcal{M}_{\text{P,R}}(l)}^- \sum_{(\mathbf{v},\mathbf{w})\in M_{P,G}^-(l,t)} \sum_{m=0}^{M} f_m(\mathbf{v},\mathbf{w}),$$
$$\forall l\in\mathcal{L}, t\in\{1,\ldots,T\}. \tag{7}$$

We are now in a position to state the Power-in-the-loop AMoD (P-AMoD) problem:

$$\underset{f_m,\lambda_m^{c,\text{in}},\lambda_m^{c,t,\text{out}},N_F,\theta,p}{\text{minimize}} \quad V_T T_M + V_D D_v + \sum_{t=1}^{T}\sum_{g\in\mathcal{G}} o_g(t)p(g,t), \tag{8}$$
$$\text{subject to} \qquad (1),(2),(3),(4),(6), \text{ and } (7).$$

D. Discussion

Some comments are in order. First, the model assumes that the TSO and the ISO share the goal of maximizing social welfare and are willing to collaborate on a joint policy. This assumption is, in general, not realistic: not only do the TSO and ISO have different goals, but they are also generally reluctant to share the information required for successful coordination. However, once a socially optimal strategy is found, efficient coordination mechanisms can be designed that steer rational agents towards that strategy. In the extended version of this paper [19] we show that the social optimum can be enforced as a general equilibrium through appropriate pricing, and that the market-clearing prices can be computed with a privacy-preserving mechanism that does not require the agents to disclose their private information.

Second, we consider single-occupancy vehicles, in line with the mode of operations of current MoD systems; the extension to ride-sharing is an interesting avenue for future research.

Third, the network flow model has some well-known limitations: chiefly, it does not capture the stochasticity of the customer arrival process, and it does not directly yield *integral* routes suitable for real-time control of vehicles. Furthermore, in this paper, customer requests are assumed to be deterministic and known in advance, an assumption that is not consistent with the paradigm of on-demand mobility. To overcome these limitation, in Section V we propose a receding-horizon implementation of Problem 8. Moreover, future requests may be interpreted as the *expected* number of future transportation requests (which may be learned from historical data and

demand models): accordingly, the model proposed in this section may be used for planning on timescales of days and hours, akin to the Day-Ahead-Market already in use in the electric power network [14].

Finally, the DC model for the power network has some shortcomings, chiefly the inability to handle voltage constraints [10] and system-dependent accuracy [23]. On the other hand, its linearity makes it amenable to large-scale optimization and easy to integrate within the economic theory upon which the transmission-oriented market design is based on [23]. Moreover, the DC model is widely adopted among ISOs [16], and its LMP calculations are fairly accurate [17]. Hence, the DC model is appropriate for high-level synthesis of joint control policies such as those considered in this paper.

III. SOLUTION ALGORITHMS

The number of variables of the P-AMoD problem (8) is $(M+1)|\mathcal{E}| + MC(T+1) + |\mathcal{V}_R|C + T(|\mathcal{G}| + |\mathcal{B}|)$. The size of the edge set $\mathcal{E}$ is $|\mathcal{E}| = \Theta((|\mathcal{E}_R| + |\mathcal{S}|)CT)$, and the number of customer demands admits an upper bound $M = O(|\mathcal{V}_R|^2T)$, since each customer demand is associated with an origin, a destination, and a departure time. The size of the problem is dominated by the customer flow variables in the road network – the number of such variables is $M|\mathcal{E}| = O((|\mathcal{V}_R|^2T)(|\mathcal{E}_R| + |\mathcal{S}|)CT)$. Consider a typical problem with 25 road nodes, 200 road links, 30 charge levels, and a horizon of 20 time steps. Such problem results in a number of variables on the order of $2 \cdot 10^9$, which can not be solved even by state-of-the-art solvers on modern hardware [15].

In this section, we propose a *bundling* procedure that collects multiple customer demands in a single customer flow without loss of information. The procedure allows one to reduce the number of network flows to $O(|\mathcal{V}_R|)$: as a result, the size of the prototypical problem above is reduced to $4 \cdot 10^6$ variables, well within reach of modern solvers. The procedure relies on the notion of *bundled customer flow*,

Definition III.1 (Bundled customer flow). *Consider the set of customer requests $\{v_m, w_m, t_m, \lambda_m\}_{m=1}^M$. Denote the set of customer destinations as $\mathcal{D} := \{\cup_{m=1}^M w_m\}$. For a given destination $d_B \in \mathcal{D}$, we define a* bundled customer flow *as a function $f_{B,d_B}(\mathbf{u}, \mathbf{v}) : \mathcal{E} \mapsto \mathbb{R}_{\geq 0}$ that satisfies*

$$\sum_{\mathbf{u}:(\mathbf{u},\mathbf{v})\in\mathcal{E}} f_{B,d_B}(\mathbf{u}, \mathbf{v}) + \sum_{\substack{m\in\{1,\ldots,M\}:\\ w_m=d_B}} 1_{v_{\mathbf{v}}=v_m} 1_{t_{\mathbf{v}}=t_m} \lambda_m^{c_{\mathbf{v}},in}$$
$$= \sum_{\mathbf{w}:(\mathbf{v},\mathbf{w})\in\mathcal{E}} f_{B,d_B}(\mathbf{v}, \mathbf{w}) + 1_{v_{\mathbf{v}}=d_B} \sum_{\substack{m\in\{1,\ldots,M\}:\\ w_m=d_B}} \lambda_m^{t_{\mathbf{v}},c_{\mathbf{v}},out}, \ \forall \mathbf{v} \in \mathcal{V}, \quad (9a)$$

$$\sum_{c=1}^C \lambda_m^{c,in} = \lambda_m, \ \forall m \in \{1,\ldots,M\} : w_m = d_B, \quad (9b)$$

$$\sum_{\substack{m\in\{1,\ldots,M\}:\\ w_m=d_B}} \sum_{t=1}^T \sum_{c=1}^C \lambda_m^{t,c,out} = \sum_{\substack{m\in\{1,\ldots,M\}:\\ w_m=d_B}} \lambda_m \quad (9c)$$

Intuitively, the bundled customer flow for a given destination d_B can be thought of as the sum of customer flows (i.e., network flows satisfying Eq. (1)) for *all* customer requests whose destination is node d_B. A bundled customer flow is an *equivalent representation* for a set of customer flows belonging

to customer requests sharing the same destination. The next lemma formalizes this intuition.

Lemma III.2 (Equivalency between customer flows and bundled customer flows). *Consider a network $G(\mathcal{V}, \mathcal{E})$ and a set of customer requests $\{v_m, w_m, t_m, \lambda_m\}_{m=1}^M$. Assume there exists a bundled customer flow $\{f_{B,d_B}(\mathbf{u},\mathbf{v})\}_{(\mathbf{u},\mathbf{v})\in\mathcal{E}}$ that satisfies Equation (9) for a destination $d_B \in \mathcal{D}$. Then, for each customer request $\{v_m, d_B, t_m, \lambda_m\}$ with destination d_B, there exists a customer flow $f_m(\mathbf{u},\mathbf{v})$ that satisfies Eq. (1). Furthermore, for each edge $(\mathbf{u},\mathbf{v}) \in \mathcal{E}$, $f_{B,d_B}(\mathbf{u},\mathbf{v}) = \sum_{m\in\{1,\ldots,M\}:w_m=d_B} f_m(\mathbf{u},\mathbf{v})$.*

Proof sketch: The proof is constructive. Define as path flow a network flow that has a fixed intensity on edges belonging to a path without cycles from the origin to the destination and zero otherwise. The flow decomposition algorithm [1, Ch. 3.5] is used to decompose the bundled customer flow into a collection of path flows, each with a single origin node $\mathbf{v} \in \mathcal{V}$ and destination node $\mathbf{w} \in \mathcal{V}$ with $v_{\mathbf{w}} = d_B$. The customer flow for customer request (v_m, d_B, t, λ) is then obtained as the sum of path flows leaving origin nodes $\{\mathbf{v} = (v_m, t_m, c)\}_{c=1}^C$ with total intensity λ_m.

We can leverage the result in Lemma III.2 to solve the P-AMoD problem in terms of bundled customer flows, thus dramatically decreasing the problem size. The next theorem formalizes this intuition.

Theorem III.3 (P-AMoD with bundled customer flows). *Consider the following problem, referred to as the* bundled P-AMoD problem*:*

$$\underset{\substack{f_0, f_{B,d_B}, \lambda_m^{c,in},\\ \lambda_m^{c,t,out}, N_F, \theta, p}}{minimize} \sum_{m=1}^M V_T T_M + V_D D_v + \sum_{t=1}^T \sum_{g\in\mathcal{G}} o_g(t)p(g,t), \quad (10)$$

$$subject\ to \quad (9) \quad \forall d_B \in \mathcal{D},$$
$$(2),(3),(4),(6),\ and\ (7),$$

where each instance of $\sum_{m=1}^M f_m$ in the cost function and in Equations (2), (3), (4), (6), and (7) is replaced by $\sum_{d_B\in\mathcal{D}} f_{B,d_B}$. The bundled P-AMoD problem (10) admits a feasible solution if and only if the P-AMoD problem (8) admits a solution. Furthermore, the optimal values of Problem (8) and Problem (10) are equal.

Proof sketch: The proof follows directly from Lemma III.2 and the linearity of the cost function. We refer the reader to the extended version [19] for a rigorous proof.

The optimization problem in (10) can be solved with a number of variables that is $O((|\mathcal{V}_R|+1)|\mathcal{E}| + MC + |\mathcal{V}_R|C + T(|\mathcal{G}| + |\mathcal{E}_p| + |\mathcal{B}|))$. To see this, note that in Equation (9) the variables $\{\lambda_m^{t,c,out}\}_{\{m,t,c\}}$ only appear as part of the sum $\sum_{m\in\{1,\ldots,M\}:w_m=d_B} \lambda_m^{t,c,out}$ and therefore may be replaced by the smaller set of variables $\{\lambda_{d_B}^{t,c,out}\}_{\{d_B,t,c\}}$, where $\lambda_{d_B}^{t,c,out} := \sum_{m\in\{1,\ldots,M\}:w_m=d_B} \lambda_m^{t,c,out}$, without loss of generality. Compared to Problem (8), the number of customer flow variables. which dominate the problem size, grows linearly (as opposed to quadratically) with the number of nodes $|\mathcal{V}_R|$ and does not depend on the time horizon T.

IV. NUMERICAL EXPERIMENTS

We study a hypothetical deployment of an AMoD system to satisfy medium-distance commuting needs in the Dallas-Fort

Worth metroplex, with the primary objective of investigating the interaction between such system and the Texas power network. Specifically, we study a ten-hour interval corresponding to one commuting cycle, from 5 a.m. to 3 p.m., with 30-minute resolution. Data on commuting patterns is collected from the 2006-2010 Census Tract Flows, based on the American Communities Survey. Census tracts in the metroplex are aggregated in 25 clusters, as shown in Figure 2. We only consider trips starting and ending in different clusters: the total number of customer requests is 400,532. The commuters' value of time is set equal to $24.40/hr, in accordance with DOT guidelines [26]. The road network, the road capacities, and the travel times are obtained from OpenStreetMap data and simplified. The resulting road network, containing 25 nodes and 147 road links, is shown in Figure 2.
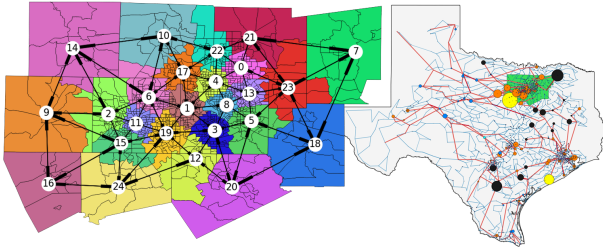

Fig. 2. Left: Census tracts and simplified road network for Dallas-Fort Worth. Right: Texas power network model (from [11]).

The battery capacity and power consumption of the EVs are modeled after the 2017 Chevrolet Bolt [5]. The cost of operation of the vehicles, excluding electricity, is 48.6c/mile, in accordance with DOT guidelines [4]. The fleet consists of 150,000 EVs, i.e. 1 AMoD vehicle for every 2.67 customers, similar to the 2.6 ratio in [22]. To represent the possibility that vehicles might not begin the day fully charged, each EV starts the day with a 50% battery charge and is required to have the same level of charge at the end of the simulation.

We adopt a synthetic model of the Texas power network provided in [11] and portrayed in Figure 2. The model provided does not contain power generation costs: we labeled each generator according to its source of power and assigned generation costs according to U.S. Energy Information Administration estimates [6]. The model is also time-invariant; to model the time evolution of power loads and the availability of solar and wind power we used historical data from ERCOT, Texas's ISO, and we imposed ramp-up and ramp-down constraints of 10%/hr and 40%/hr on the generation capability of nuclear and coal power plants, respectively.

We compare the results of three simulation studies. In the baseline simulation study, no electric vehicles are present: we consider the power network *in isolation* subject only to exogenous loads. In the P-AMoD simulation study, we solve Problem (10), which embodies the cooperation between the TSO and the ISO. Finally, in the uncoordinated simulation study, we first solve the TSO's Vehicle Routing and Charging problem with *fixed* electricity prices obtained from the baseline simulation study; we then compute the load on the power network resulting from the vehicles' charging and discharging, and solve the ISO's Economic Dispatch problem with the updated loads. The uncoordinated simulation study captures the scenario where the TSO attempts to minimize its passengers' cost while disregarding the coupling with the power network.

Table I and Figure 3 show the results.

TABLE I
SIMULATION RESULTS (ONE COMMUTING CYCLE, 10 HOURS).

	Baseline	P-AMoD	Uncoord.
Avg. cust. travel time [h]	-	1.2532	1.2532
Tot. energy demand [GWh]	517.498	520.541	520.979
Tot. electricity expenditure [k$]	39,604.71	39,264.84	39,629.50
w.r.t. baseline [k$]		-339.87	+24.79
Avg. price in DFW [$/MW]	78.68	75.79	77.47
TSO tot. elec. expenditure [k$]	-	227.98	296.82

The quality of service experienced by TSO customers, measured by the average travel time, is virtually identical in the P-AMoD and in the uncoordinated case. The energy demand of the AMoD system is also very similar in both cases. On the other hand, the effect of coordination on the overall electricity expenditure is noticeable. Coordination between the TSO and the ISO causes a *reduction* in the total expenditure for electricity of $339,870 per commuting cycle compared to the baseline case, despite the increased demand! In other words, a P-AMoD system allows a TSO to deliver on-demand transportation without an increase in overall electricity expenditure – a remarkable, and perhaps surprising, finding. Instead, in the uncoordinated case, the total expenditure for electricity is *increased* by $24,790. This corresponds to a difference of $364,066 between the P-AMoD case and the uncoordinated case, which compounds to savings in electricity expenditure of $182M per year (assuming two commuting cycles per day and 250 work days per year).

Who benefits from the reduction in energy expenditure? From the last two rows in Table I, one can see that the average price of electricity in the P-AMoD case is 2.16% lower than in the uncoordinated case in Dallas-Fort Worth (corresponding to savings of $122.3M/year). The energy expenditure of the TSO in the P-AMoD case is 23.5% lower than in the uncoordinated case (a saving of $69,740 per commuting cycle, corresponding to close to $35M/year). Finally, electricity customers outside of Dallas experience a small reduction of 0.75% in their energy expenditure. Thus, the majority of the benefits of coordination are reaped by customers of the power network in the region where the AMoD system is deployed; the TSO also benefits from a noticeable reduction in its electricity expenditure.
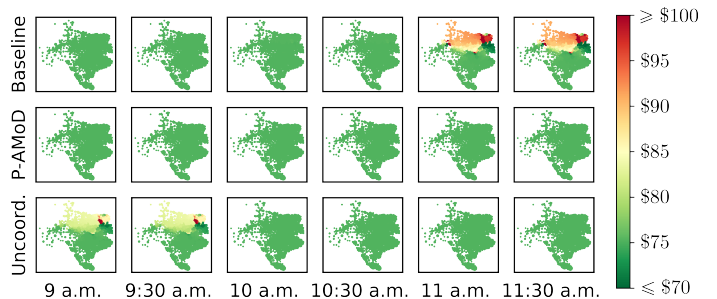

Fig. 3. LMPs in Texas between 9 a.m. and 11:30 a.m. The presence of the AMoD fleet can reduce locational marginal prices; coordination between the TSO and the ISO can yield a further reduction.

Figure 3 shows this phenomenon in detail. The presence of the AMoD system results in a decrease in the LMPs with respect to the baseline case (11-11:30 a.m.). As electricity prices increase, empty vehicles travel to carefully chosen stations to sell their stored energy back to the network: this results in reduced congestion and lower prices in the power network,

even in the absence of coordination. Crucially, coordination between the TSO and the ISO can results in further decreases in the price of electricity with respect to the uncoordinated case (9-9:30 a.m.), significantly curtailing the impact of the AMoD system on the power network. By leveraging their battery capacities and acting as mobile storage units, the EVs are able to reduce congestion in the power transmission network: this results in lower LMPs in the Dallas-Fort Worth region, and hence lower electricity expenditure. Simulations were carried out on commodity hardware (Intel Core i7-5960, 64 GB RAM) and used the MOSEK LP solver. The simulations required 3,923s for the P-AMoD scenario, 2,885s for the uncoordinated scenario, and 4.55s for the baseline scenario. While such computation times could be improved by using high-performance computational hardware, in the next section we present a receding-horizon algorithm for P-AMoD which, in addition to the intrinsic robustness benefits of closed-loop control, can be solved in seconds on commodity hardware.

V. A Receding-Horizon Algorithm for P-AMoD

Leveraging the structural insights from the network flow optimization problem of the previous sections, along with a few mild assumptions, we next devise a receding-horizon algorithm that is robust to the uncertainty in future demand. Additionally, this algorithm trades off some suboptimality, which we characterize with simulations, for very fast computation times. Due to space limitations, we only provide a high-level description of the algorithm: a detailed description will be provided in the journal version of this paper.

To reduce the computational complexity of the optimization problem, we *decouple* the customer routing process from the P-AMoD optimization. The key assumption is that customer-carrying trips follow pre-computed routes and are never interrupted by a charging/discharging event. Formally, customer trips from node $i \in \mathcal{V}_R$ to node $j \in \mathcal{V}_R$ follow a fixed route with a travel time of $t_{i \to j}$ and a required charge of $c_{i \to j}$. Thus, customer flows $\{f_{B,d_B}(\mathbf{u}, \mathbf{v})\}_{(\mathbf{u},\mathbf{v}),d_B}$ are no longer part of the optimization variables and Equation (9a) is redundant. However, the initial and final charge of the customer-carrying vehicles $\{\lambda_m^{c,\text{in}}\}$ and $\{\lambda_m^{t,c,\text{out}}\}$ remain optimization variables. The following constraint ensures that charge is conserved along customer routes, that is, that vehicles traveling from i to j and departing at time t at charge level c arrive at time $t + t_{i \to j}$ with charge $c - c_{i \to j}$:

$$\lambda_m^{t,c,\text{out}} = \begin{cases} \lambda_m^{c+c_{v_m \to w_m},\text{in}} & \text{if } t_m = t - t_{v_m \to w_m} \\ 0 & \text{otherwise} \end{cases} \quad (11)$$

for all $t \in \{1, \ldots, T\}, c \in \{1, \ldots, C\}, m \in \{1, \ldots, M\}$. The cost function is also modified to remove the customers' travel times, and road congestion constraints are adjusted to account for the traffic induced by customer-carrying vehicles.

The optimization problem is solved in receding-horizon fashion. The solution to the problem is, in general, fractional: control actions for the vehicles are computed by *sampling* the first time step of the solution.

We assess the performance of the receding-horizon P-AMoD controller with an agent-based simulation based on the same case study considered in the previous section. The receding-horizon problem is solved every two minutes with a 2-hour lookahead and a 30-minute time step. The performance

of the algorithm is compared with an uncoordinated receding-horizon controller that optimizes the AMoD system's operations under the assumption that electricity prices stay constant.

Coordination results in savings of approximately $150,000 per commuting cycle (corresponding to almost $75M/year) with respect to the uncoordinated algorithm. In particular, cooperation results in 0.88% lower electricity prices in Dallas-Fort Worth. The TSO's expense in the coordinated case is comparable with the expense computed in Section IV, which represents an *upper bound* on the performance of a receding-horizon P-AMoD controller. On the other hand, the average price of electricity in Texas increases by 1.7% compared to the baseline case studied in Section IV. This is not unexpected, as we use a short 2-hour lookahead. An important direction of future research is to perform a detailed sensitivity analysis, and in particular to explore how the lookahead time affects the tradeoff between computational complexity, economic savings, and robustness to inaccuracies in demand forecasting. The receding-horizon P-AMoD problem was solved in an average of 2.09s; thus, the algorithm is amenable to closed-loop control of large-scale systems.

VI. Conclusions and Future Work

In this paper we studied the interaction between an AMoD system and the power network. The model we proposed subsumes earlier models for AMoD systems and for the power network; critically, it captures the coupling between the two systems and allows for their *joint optimization*. We also proposed a numerical procedure to losslessly reduce the dimensionality of the P-AMoD optimization problem (Eq. (10)), making realistic problems amenable to efficient numerical optimization on commodity hardware. We applied our model to a case study of an AMoD deployment in Dallas-Fort Worth, TX. The case study showed that coordination between the TSO and the ISO can result in a *reduction* in the overall electricity expenditure (despite the increase in demand), while having a negligible impact on the TSO's quality of service. Finally, we presented a receding-horizon algorithm for P-AMoD that delivers computation times in the order of seconds and provides built-in robustness to uncertainty in future demand at the price of some suboptimality.

This work opens multiple avenues of research. First, the model in this paper assumes that the TSO and ISO are willing to collaborate and share their private information. In the extended version of this paper [19], we propose a pricing scheme that enforces the social optimum as a general equilibrium and a distributed algorithm that allows the agents to compute the market-clearing prices without sharing private information. Second, we plan to capture the impact of cooperation between the TSO and the ISO on the power *distribution* network by incorporating convex optimal power flow models. Third, the model of the power network considered in this paper does not capture ancillary services such as regulation and spinning reserves. We will extend our model to capture those and evaluate the feasibility of using coordinated fleets of EVs to aid in short-term control of the power network. Finally, we wish to explore the effect of TSO-ISO coordination on penetration of renewable energy sources, and to determine whether large-scale deployment of AMoD systems can increase the fraction of renewable power sources in the generation power mix.

REFERENCES

[1] R. K. Ahuja, T. L. Magnanti, and J. B. Orlin. *Network Flows: Theory, Algorithms and Applications*. Prentice Hall, 1993.

[2] M. Alizadeh, H.-T. Wai, A. Scaglione, A. Goldsmith, Y. Y. Fan, and T. Javidi. Optimized path planning for electric vehicle routing and charging. In *Allerton Conf. on Communications, Control and Computing*, 2014.

[3] M. Alizadeh, H.-T. Wai, M. Chowdhury, A. Goldsmith, A. Scaglione, and T. Javidi. Optimal pricing to manage electric vehicles in coupled power and transportation networks. *IEEE Transactions on Control of Network Systems*, 4(4):863–875, 2017.

[4] Bureau of Transportation Statistics. National transportation statistics. Technical report, U.S. Dept. of Transportation, 2016.

[5] Chevrolet. Bolt EV. Available at http://www.chevrolet.com/bolt-ev-electric-vehicle. Retrieved on June 5, 2017., 2017.

[6] EIA. Levelized cost and levelized avoided cost of new generation resources in the annual energy outlook 2017. Technical report, U.S. Energy Information Administration, 2017.

[7] J.D. Glover, M.S. Sarma, and T. Overbye. *Power System Analysis and Design*. Cengage Learning, fifth edition, 2011.

[8] D. Goeke and M. Schneider. Routing a mixed fleet of electric and conventional vehicles. *European Journal of Operational Research*, 245(1):81–99, 2015.

[9] S. W. Hadley and A. A. Tsvetkova. Potential impacts of plug-in hybrid electric vehicles on regional power generation. *The Electricity Journal*, 22(10):56–68, 2009.

[10] W. W. Hogan. Markets in real electric networks require reactive prices. In *Electricity Transmission Pricing and Technology*, chapter 7. Springer Netherlands, Dordrecht, 1996.

[11] Illinois Center for a Smarter Electric Grid (ICSEG). Texas 2000-June 2016 synthetic power case, 2016. URL http://icseg.iti.illinois.edu/synthetic-power-cases/texas2000-june2016/.

[12] W. Kempton and J. Tomić. Vehicle-to-grid power fundamentals: Calculating capacity and net revenue. *Journal of Power Sources*, 144(1):268–279, 2005.

[13] M. E. Khodayar, L. Wu, and Z. Li. Electric vehicle mobility in transmission-constrained hourly power generation scheduling. *IEEE Transactions on Smart Grid*, 4(2):779–788, 2013.

[14] D. S. Kirschen and G. Strbac. *Fundamentals of Power System Economics*. John Wiley & Sons, first edition, 2004.

[15] Hans D. Mittelmann. Decision tree for optimization software. http://plato.asu.edu/guide, 2016. URL http://plato.asu.edu/guide.

[16] R. P. O'Neill, T. Dautel, and E. Krall. Recent ISO software enhancements and future software and modeling plans. Technical report, Federal Energy Regulatory Commission, 2011.

[17] T. J. Overbye, X. Cheng, and Y. Sun. A comparison of the AC and DC power flow models for LMP calculations. In *Hawaii Int. Conf. on System Sciences*, 2004.

[18] S. Pourazarm, C. G. Cassandras, and T. Wang. Optimal routing and charging of energy-limited vehicles in traffic networks. *Int. Journal of Robust and Nonlinear Control*, 26(6):1325–1350, 2016.

[19] F. Rossi, R. Iglesias, M. Alizadeh, and M. Pavone. On the interaction between Autonomous Mobility-on-Demand systems and the power network: Models and coordination algorithms. In *Robotics: Science and Systems*, 2018. Extended version available at https://arxiv.org/abs/1709.04906.

[20] N. Rotering and M. Ilic. Optimal charge control of plug-in hybrid electric vehicles in deregulated electricity markets. *IEEE Transactions on Power Systems*, 26(3):1021–1029, 2011.

[21] R. Sioshansi. OR Forum—modeling the impacts of electricity tariffs on plug-in hybrid electric vehicle charging, costs, and emissions. *Operations Research*, 60(3):506–516, 2012.

[22] K. Spieser, K. Treleaven, R. Zhang, E. Frazzoli, D. Morton, and M. Pavone. Toward a systematic approach to the design and evaluation of Autonomous Mobility-on-Demand systems: A case study in Singapore. In *Road Vehicle Automation*. Springer, 2014.

[23] B. Stott, J. Jardim, and O. Alsaç. DC power flow revisited. *IEEE Transactions on Power Systems*, 24(3):1290–1300, 2009.

[24] K. Turitsyn, N. Sinitsyn, S. Backhaus, and M. Chertkov. Robust broadcast-communication control of electric vehicle charging. In *IEEE Int. Conf. on Smart Grid Communications (SmartGridComm)*, 2010.

[25] W. Tushar, W. Saad, H. V. Poor, and D. B. Smith. Economics of electric vehicle charging: A game theoretic approach. *IEEE Transactions on Power Systems*, 3(4):1767–1778, 2012.

[26] U.S. Dept. of Transportation. Revised departmental guidance on valuation of travel time in economic analysis. Technical report, 2015.

[27] L. Wang, A. Lin, and Y. Chen. Potential impact of recharging plug-in hybrid electric vehicles on locational marginal prices. *Naval Research Logistics*, 57(8):686–700, 2010.

[28] J. G. Wardrop. Some theoretical aspects of road traffic research. *Proc. of the Institution of Civil Engineers*, 1(3):325–362, 1952.

[29] World Health Organization. 7 million premature deaths annually linked to air pollution. http://www.who.int/mediacentre/news/releases/2014/air-pollution/en/. Retrieved on March 2, 2018., 2014.

[30] R. Zhang, F. Rossi, and M. Pavone. Routing autonomous vehicles in congested transportation networks: Structural properties and coordination algorithms. In *Robotics: Science and Systems*, 2016.