

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—We study 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* model that captures the coupling between the two systems stemming from the vehicles’ charging requirements. The model captures time-varying customer demand and power generation costs, road congestion, battery depreciation, 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. Next, we show that the socially optimal solution to the joint problem can be enforced as a general equilibrium, and we provide a dual decomposition algorithm that allows the transportation and power network operators to compute the market clearing prices without sharing private information. We assess the performance of the model and algorithms by studying a hypothetical electric-powered AMoD system in Dallas-Fort Worth and its impact on the Texas power network. Lack of coordination between the AMoD system and the power network would cause a 4.4% increase in the price of electricity in Dallas-Fort Worth; conversely, depending on the maturity of battery technology, coordination between the AMoD system and the power network would *reduce* the total electricity expenditure compared to the case where no cars are present (despite the increased demand for electricity) and yield savings of up to \$147M/year for local power network customers. Agent-based simulations with receding-horizon versions of the algorithms further corroborate these findings. 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 [1]. 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.

Federico Rossi is with the the Jet Propulsion Laboratory, California Institute of Technology, Pasadena (CA) 91109. Email: federico.rossi@jpl.nasa.gov.

Marco Pavone is with the Department of Aeronautics and Astronautics, Stanford University, Stanford (CA) 94305. Email: pavone@stanford.edu.

Ramon Iglesias is with the Department of Civil and Environmental Engineering, Stanford University, Stanford (CA) 94305. Email: rdit@stanford.edu.

Mahnoosh Alizadeh is with the Electrical & Computer Engineering Department, University of California, Santa Barbara, Santa Barbara, CA 93106. Email: alizadeh@ucsb.edu.

Federico Rossi worked on this paper while he was a Ph.D. student at Stanford University. This research was supported by the National Science Foundation under CAREER Award CMMI-1454737 and by the Toyota Research Institute (TRI). This article solely reflects the opinions and conclusions of its authors and not NSF, TRI or any other Toyota entity.

* Corresponding author.

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 [2] 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) [3], [4], [5]. For example, [5] 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. [4] 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, lower electricity costs, and increased integration of renewable energy sources). 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: Control of AMoD systems has been

addressed in multiple lines of work, including queuing-theoretical approaches [6], [7], network flow approaches [8], [9], integer linear programming and model-predictive control approaches [10], [11], and simulation-based approaches [12], [13]. However, throughout these works, AMoD systems are assumed to have no impact on the electric power network.

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 [14], [15], [16]. This assumption is also commonly made when selecting the locations of charging stations (i.e., optimizing the charging profile in *space*) [17], [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 [5] and game-theoretical [3], [19] settings. Accordingly, [4] 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 [20], where it is shown that widespread adoption of EVs and V2G could foster significantly increased adoption of wind power. Going one step further, [21] 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, [20] does not capture the *spatial* component of the power and transportation networks, while [21] assumes that the vehicles' schedules are fixed.

The objective of this paper is to investigate the interaction between AMoD and the electric power network (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. Our work improves upon the state of the art (in particular, [4]) along three main dimensions: (i) it provides rigorous models for a fleet of *shared* and *autonomous* EVs; (ii) it provides efficient algorithms that can scale to large-scale instances; and (iii) it characterizes the vehicles' ability to return power to the power network through vehicle-to-grid (V2G) schemes, and its economic benefits.

Statement of contributions: First, we propose a joint model for P-AMoD systems. 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 electricity generation costs, congestion in the road network, vehicle battery depreciation, 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 optimized on commodity hardware. Third,

we show that the socially-optimal solution to the P-AMoD problem can be enforced as a general equilibrium through pricing, and we propose a distributed privacy-preserving algorithm that the transportation and power network operators can employ to compute the efficient (market clearing) prices without disclosing their private information. Fourth, 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. Agent-based simulations with receding-horizon versions of the algorithms show that, in absence of coordination, large-scale adoption of electric AMoD can cause widespread blackouts and increase electricity prices by almost 50%; conversely, the receding-horizon P-AMoD algorithm is able to maintain electricity prices constant, despite the substantial increase in power demand.

A preliminary version of this paper was presented at the 2018 Robotics: Science and Systems conference. In this revised and extended version, we provide as additional contributions (i) a rigorous proof that the socially-optimal solution can be enforced as a general equilibrium, (ii) a privacy-preserving distributed optimization algorithm, (iii) additional numerical results, and (iv) proofs of all theorems.

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 show that the socially optimal solution to the P-AMoD problem can be enforced as a general equilibrium and propose a privacy-preserving distributed optimization algorithm. In Section V, we evaluate our model and algorithm on a case study of Dallas-Fort Worth. In Section VI, we draw conclusions and discuss directions for future work. Finally, in the Appendix, we present agent-based simulations and proofs of all theorems.

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 [9] to a time-varying, charge-aware 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 network 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 used to calculate market clearing prices for electricity. The power transmission network comprises spatially-distributed energy providers that are connected to spatially-distributed power network users 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: these systems induce 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 power demands 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 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 $\bar{f}_{v,w} \in \mathbb{R}_{\geq 0}$ captures the maximum vehicle flow rate (i.e., the number of vehicles per unit of time) that the road link can accommodate on top of exogenous traffic 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 $\mathcal{S} \subset \mathcal{V}_R$. Each charging station $s \in \mathcal{S}$ is characterized by a charging rate $\delta c_s^+ \in \{1, \dots, C\}$, a discharging rate $\delta c_s^- \in \{-C, \dots, -1\}$, a time-varying charging price $p_s^+(t) \in \mathbb{R}$, a time-varying discharging price $p_s^-(t) \in \mathbb{R}$, and vehicle capacity $\bar{S}_s \in \mathbb{N}$. 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 $p_s^+(t)$ and $p_s^-(t)$ capture the cost of one discrete unit charge level (or, conversely, the payment the vehicles receive for returning one unit charge level to the grid) at time t ; in this paper, we assume that $p_s^+(t) = p_s^-(t)$ (in accordance with the assumption of an arbitrage-free market). The vehicle capacity \bar{S}_s models the maximum number of vehicles that can simultaneously charge or discharge at station s . Charging and discharging (due both to driving activity and to vehicle-to-grid power injection) cause wear in the vehicles’ batteries. The battery depreciation per unit charge or discharge is denoted as d_B . Battery depreciation captures the cost of replacing a battery at the end of its useful life; note, however, that the vehicle’s battery capacity is assumed to remain constant during the model’s finite horizon.

Expanded AMoD network: We are now in a position to rigorously define the network flow model for the AMoD system. We introduce an *expanded* AMoD network modeled as a directed graph $G = (\mathcal{V}, \mathcal{E})$. The graph G captures the time-varying nature of the problem and tracks the battery charge level of the autonomous vehicles. Specifically, nodes $\mathbf{v} \in \mathcal{V}$ model physical locations 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 subsets, namely \mathcal{E}_L and \mathcal{E}_S , such that $\mathcal{E}_L \cup \mathcal{E}_S = \mathcal{E}$ and $\mathcal{E}_L \cap \mathcal{E}_S = \emptyset$. 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 .

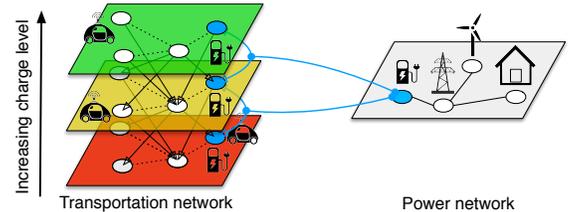


Fig. 1. Augmented transportation and power networks. As vehicles travel on road links (modeled by solid black arrows), their charge level decreases. Blue nodes represent charging stations: the flows on charging and discharging edges affect the load at the corresponding nodes in the power network. For 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 within 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 [22]. 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 [22, 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_v to location v_w (or charging/discharging at location $v_v = v_w$) from time t_v to time t_w , with an initial battery charge of c_v and a final battery charge of c_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:

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

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

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 are fixed; the final position and charge level are optimization variables (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 $t_v = 0$, denotes the number of rebalancing vehicles entering the AMoD system at location v_v at time t_v with charge level c_v . Conversely, $N_F(\mathbf{v})$, with $t_v = T$ denotes the number of rebalancing vehicles at location v_v at time t_v with charge level c_v . For $t_v \neq 0$, $N_I(\mathbf{v}) = 0$; for $t_v \neq T$, $N_F(\mathbf{v}) = 0$. 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:

$$\sum_{\mathbf{u}: (\mathbf{u}, \mathbf{v}) \in \mathcal{E}} f_0(\mathbf{u}, \mathbf{v}) + \sum_{m=1}^M 1_{v_v = w_m} \lambda_m^{t_v, c_v, \text{out}} + N_I(\mathbf{v}) = \sum_{\mathbf{w}: (\mathbf{v}, \mathbf{w}) \in \mathcal{E}} f_0(\mathbf{v}, \mathbf{w}) + \sum_{m=1}^M 1_{v_v = v_m} 1_{t_v = t_m} \lambda_m^{c_v, \text{in}} + N_F(\mathbf{v}), \quad \forall \mathbf{v} \in \mathcal{V}. \quad (2)$$

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 [9] and is consistent with classical traffic flow theory [23]. 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:

$$\sum_{c_v=1}^C \sum_{m=0}^M f_m(\mathbf{v}, \mathbf{w}) \leq \bar{f}_{(v_v, v_w)}, \quad \forall (v_v, v_w) \in \mathcal{E}_R, t_v \in \{1, \dots, T\}. \quad (3)$$

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

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

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

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

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

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, $d_{v_v, v_w} = 0$. The total cost of electricity incurred by the vehicles (including any credit from selling electricity to the power network) is

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

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

The overall battery depreciation due to charging and discharging is

$$V_B = d_B \sum_{m=0}^M \left[\sum_{(\mathbf{v}, \mathbf{w}) \in \mathcal{E}_S} f_m(\mathbf{v}, \mathbf{w}) |\delta c_{v, w}| + \sum_{(\mathbf{v}, \mathbf{w}) \in \mathcal{E}_L} f_m(\mathbf{v}, \mathbf{w}) |c_{(v_v, v_w)}| \right].$$

(Note that battery depreciation accounts for both charging and discharging, since battery life is determined by the number of charging/discharging cycles incurred by the battery cells).

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 time interval) as V_T and the operation cost per kilometer of the vehicles

(excluding electricity costs) as V_D . We are now in a position to state the TSO's Vehicle Routing and Charging problem:

$$\begin{aligned} & \underset{f_m, \lambda_m^{c, \text{in}}, \lambda_m^{t, c, \text{out}}, N_F}{\text{minimize}} && V_D D_V + V_E + V_B + V_T T_M, && (5a) \\ & \text{subject to} && (1), (2), (3), \text{ and } (4). && (5b) \end{aligned}$$

B. Linear model of power network

In this paper, the power network is modeled according to the well-known DC model [24, 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 [25]. 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. 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 $\bar{p}_g(t)$, 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 \bar{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 substation transformers are modeled by an upper bound $\bar{d}_l(t)$ on the power that can be delivered at each load node.

We define a generator power function $p : (\mathcal{G}, \{1, \dots, T\}) \mapsto \mathbb{R}_{\geq 0}$, and a phase angle function $\theta : (\mathcal{B}, \{1, \dots, T\}) \mapsto \mathbb{R}$. The generation cost is defined as

$$C_G = \sum_{t=1}^T \sum_{g \in \mathcal{G}} o_g(t) p(g, t).$$

The Economic Dispatch problem entails minimizing the generation cost subject to a set of feasibility constraints [24]:

$$\underset{p, \theta}{\text{minimize}} \quad C_G, \quad (6a)$$

$$\text{subject to} \quad \sum_{(u,v) \in \mathcal{E}_P} \frac{\theta(u,t) - \theta(v,t)}{x_{u,v}} + 1_{v \in \mathcal{G}} p(v,t) = 1_{v \in \mathcal{L}} d_v(t) + \sum_{(v,w) \in \mathcal{E}_P} \frac{\theta(v,t) - \theta(w,t)}{x_{v,w}}, \quad \forall v \in \mathcal{B}, t \in \{1, \dots, T\}, \quad (6b)$$

$$-\bar{p}_{b_1, b_2} \leq \frac{\theta(b_1, t) - \theta(b_2, t)}{x_{b_1, b_2}} \leq \bar{p}_{b_1, b_2}, \quad \forall (b_1, b_2) \in \mathcal{E}_P, t \in \{1, \dots, T\}, \quad (6c)$$

$$\underline{p}_g(t) \leq p(g, t) \leq \bar{p}_g(t), \quad \forall g \in \mathcal{G}, t \in \{1, \dots, T\}, \quad (6d)$$

$$-p_g^-(t) \leq p(g, t+1) - p(g, t) \leq p_g^+(t), \quad \forall g \in \mathcal{G}, t \in \{1, \dots, T-1\}, \quad (6e)$$

$$d_l(t) \leq \bar{d}_l(t), \quad \forall l \in \mathcal{L}, t \in \{1, \dots, T\}. \quad (6f)$$

Equation (6b) enforces power balance at each bus based on the so-called DC power flow equations; Equation (6c) encodes the transmission lines' thermal constraints; Equation (6d) encodes the generation capacity constraints; Equation (6e)

encodes the ramp-up and ramp-down constraints; and Equation (6f) encodes the thermal constraints of substation transformers.

The unit price of electricity at the load nodes is determined through a mechanism known as Locational Marginal Pricing (LMP) [24]. 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. Accordingly, in this paper, the LMP at each bus equals the sum of the dual variables (i.e., the shadow prices) corresponding to the power injection constraint (6b) and the substation transformer thermal constraint (6f) 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 [4].

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), pricing mechanisms can be designed to steer the system towards the optimum: we propose one such mechanism in Section IV.

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 $\mathcal{M}_{\text{PR}} : \mathcal{L} \mapsto \{\mathcal{V}_R \cup \emptyset\}$. 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, $\mathcal{M}_{\text{PG}}^+ : (\mathcal{L}, \{1, \dots, T\}) \mapsto \{\mathcal{E}_S \cup \emptyset\}$ and $\mathcal{M}_{\text{PG}}^- : (\mathcal{L}, \{1, \dots, T\}) \mapsto \{\mathcal{E}_S \cup \emptyset\}$. The function $\mathcal{M}_{\text{PG}}^+$ (respectively, $\mathcal{M}_{\text{PG}}^-$) maps a load node l and a time t to the set of charge (respectively discharge) edges in G corresponding to station $\mathcal{M}_{\text{PR}}(l)$ at time t . Formally,

$$\mathcal{M}_{\text{PG}}^+(l, t) : \{(\mathbf{v}, \mathbf{w}) \in \mathcal{E}_S \mid v_{\mathbf{v}} = v_{\mathbf{w}}, v_{\mathbf{v}} \in \mathcal{M}_{\text{PR}}(l), c_{\mathbf{v}} < c_{\mathbf{w}}, t_{\mathbf{v}} \leq t < t_{\mathbf{w}}\},$$

$$\mathcal{M}_{\text{PG}}^-(l, t) : \{(\mathbf{v}, \mathbf{w}) \in \mathcal{E}_S \mid v_{\mathbf{v}} = v_{\mathbf{w}}, v_{\mathbf{v}} \in \mathcal{M}_{\text{PR}}(l), c_{\mathbf{v}} > c_{\mathbf{w}}, t_{\mathbf{v}} \leq t < t_{\mathbf{w}}\}.$$

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,

$$\begin{aligned} d_l(t) = & d_{l,e}(t) + \frac{J_C \delta c_{\mathcal{M}_{\text{PR}}^+}(t)}{T_S \delta t_{\mathcal{M}_{\text{PR}}^+}(t)} \sum_{(\mathbf{v}, \mathbf{w}) \in \mathcal{M}_{\text{PG}}^+(l, t)} \sum_{m=0}^M f_m(\mathbf{v}, \mathbf{w}) \\ & + \frac{J_C \delta c_{\mathcal{M}_{\text{PR}}^-}(t)}{T_S \delta t_{\mathcal{M}_{\text{PR}}^-}(t)} \sum_{(\mathbf{v}, \mathbf{w}) \in \mathcal{M}_{\text{PG}}^-(l, t)} \sum_{m=0}^M f_m(\mathbf{v}, \mathbf{w}), \quad (7) \end{aligned}$$

for all $l \in \mathcal{L}, t \in \{1, \dots, T\}$.

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

$$\begin{aligned} & \underset{f_m, \lambda_m^{c,in}, \lambda_m^{t,c,out}, N_F, \theta, p}{\text{minimize}} && V_T T_M + V_D D_v + V_B + C_G, && (8a) \\ & \text{subject to} && (1), (2), (3), (4), (6), \text{ and } (7). && (8b) \end{aligned}$$

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 Section IV, we show that the social optimum can be enforced as a general equilibrium for a self-interested TSO, self-interested power generators, and a non-profit ISO acting as a market broker, and we propose a distributed *privacy-preserving* mechanism that an ISO and a TSO can adopt to compute the market-clearing prices that enforce such an equilibrium.

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. Indeed, transportation requests in our model can be interpreted as *expected* values of the corresponding stochastic processes (which can be estimated from historical data and/or via 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 [24]. Additionally, to enable real-time operations, in Appendix A we propose a receding-horizon implementation of Problem (8) that is able to quickly return integral solutions amenable to real-time control of P-AMoD systems.

Finally, the DC model for the power network has some shortcomings, chiefly the inability to handle voltage constraints [26] and system-dependent accuracy [27]. 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 [27]. Moreover, the DC model is widely adopted among ISOs [28], and its LMP calculations are fairly accurate [29]. Hence, the DC model is appropriate for high-level synthesis of joint control policies such as those considered in this paper. We remark that any convex optimal power flow model could be readily used in lieu of the DC model, since convex models are also amenable to efficient optimization and can be used to compute locational marginal prices; the study and integration of such models is an interesting direction for future research.

III. SOLUTION ALGORITHMS

The number of optimization variables in 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 $|\mathcal{E}| = \Theta((|\mathcal{E}_R| + |\mathcal{S}|)CT)$

(that is, the asymptotic growth of $|\mathcal{E}|$ is bounded both from above by a function $\bar{k}(|\mathcal{E}_R| + |\mathcal{S}|)CT$ and bounded below by a function $\underline{k}(|\mathcal{E}_R| + |\mathcal{S}|)CT$, where \bar{k} and \underline{k} are positive constants), and the number of customer requests M admits an upper bound $O(|\mathcal{V}_R|^2 T)$, 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 $M|\mathcal{E}| = O((|\mathcal{V}_R|^2 T)(|\mathcal{E}_R| + |\mathcal{S}|)CT)$. Consider a typical problem with 25 road nodes, 200 road links, 30 charge levels, and a horizon of 20 time steps. Such a 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 [30].

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

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

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

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 Equation (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 Equation (1). Furthermore, for each edge $(\mathbf{u}, \mathbf{v}) \in \mathcal{E}$, $f_{B,d_B}(\mathbf{u}, \mathbf{v}) = \sum_{m \in \{1, \dots, M\} : w_m = d_B} f_m(\mathbf{u}, \mathbf{v})$.*

Proof sketch: 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 [22, 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 $w_{\mathbf{w}} = d_B$. 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 . A rigorous proof is reported in the Appendix.

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:*

$$\begin{aligned} & \underset{f_0, f_B, d_B, \lambda_m^{c, in}, \lambda_m^{t, c, out}, N_F, \theta, p}{\text{minimize}} && V_T T_M + V_D D_v + V_B + C_G, && (10) \\ & \text{subject to} && (9) \quad \forall d_B \in \mathcal{D}, (2), (3), (4), (6), \text{ and } (7), \end{aligned}$$

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 feasible solution. Furthermore, the optimal values of Problem (8) and Problem (10) are equal.

Proof sketch: The proof follows directly from Lemma III.2. A rigorous proof can be found in the Appendix.

The optimization problem in (10) can be solved with a number of variables on the order of $O((|\mathcal{V}_R| + 1)|\mathcal{E}| + MC + |\mathcal{V}_R|C + T(|\mathcal{G}| + |\mathcal{E}_p| + |\mathcal{B}|))$. 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 $\sum_{m \in \{1, \dots, M\}: w_m = d_B} \lambda_m^{t, c, out}$ and therefore may be replaced by the smaller set of variables $\{\lambda_{d_B}^{t, c, out}\}_{\{d_B, t, c\}}$, where $\lambda_{d_B}^{t, c, out} := \sum_{m \in \{1, \dots, M\}: w_m = d_B} \lambda_m^{t, c, out}$, 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. DISTRIBUTED SOLUTION TO THE P-AMOD PROBLEM

The model and solution algorithms presented in the previous sections assume that the TSO and the ISO both wish to maximize social welfare for given generation costs; also, in order to compute the socially optimal solution to the P-AMoD problem, the TSO and the ISO must be willing to share their private information (e.g., customer transportation requests and power generation costs). In this section, we provide theoretical results and algorithmic tools to overcome these rather unrealistic assumptions. In particular, we define a P-AMoD market as a perfectly competitive market where self-interested power generators sell power to the power network, a self-interested TSO buys from and sells power to the power network, and a non-profit ISO acts as a market broker (similar to the model in [31]). In this framework, we show that the socially optimal solution to the P-AMoD problem can also be enforced as a general equilibrium [24] for the TSO and the generators in the P-AMoD market if the ISO sets the price of electricity through Locational Marginal Prices. Next, we propose a distributed privacy-preserving algorithm that the TSO and the ISO can use to compute such prices without sharing private information on transportation demand or generation costs.

A. The socially optimal solution can be enforced as a general equilibrium

Theorem IV.1 (The socially optimal solution of the P-AMoD problem can be enforced as a general equilibrium through prices). *Consider an optimal solution*

$\{f_m^, \lambda_m^{c, in*}, \lambda_m^{t, c, out*}, N_F^*, \theta^*, p^*\}$ to the P-AMoD Problem (8). Also consider a perfectly competitive market (denoted as the P-AMoD market) where a self-interested TSO solves the Vehicle Routing and Charging problem (5) by selecting variables $\{f_m, \lambda_m^{c, in}, \lambda_m^{t, c, out}, N_F\}$, self-interested power generators sell power to the network by determining the revenue-maximizing power generation schedule $\{p\}$, and a non-profit ISO acts as a market broker by setting locational marginal prices and controlling phase angles $\{\theta\}$. Then $(\{f_m^*, \lambda_m^{c, in*}, \lambda_m^{t, c, out*}, N_F^*\}, \{\theta^*\}, \{p^*\})$ is a general equilibrium.*

Proof Sketch: The proof relies on showing that satisfaction of the KKT conditions for Problem (8) implies satisfaction of the KKT conditions for Problem (5). The key insight is that the term V_E in the cost function of the Vehicle Routing and Charging problem (5) captures the marginal cost imposed by the TSO on the power network, aligning the TSO's incentives with the social optimum. A rigorous proof is reported in the Appendix.

B. A distributed algorithm for the P-AMoD problem

Next, we show that the TSO and the ISO can compute the locational marginal prices that enforce the general equilibrium without disclosing their private information. Our approach is similar to the one in [4] and relies on using a dual decomposition algorithm [32, Ch. 6.4] to solve Problem (8) in a distributed manner. Concretely, the TSO repeatedly solves Problem (5) with electricity prices proposed by the ISO, and the ISO updates the electricity prices according to the TSO's proposed charging schedule; the procedure is repeated until convergence. The TSO and the ISO only exchange publicly-available information (namely, the proposed charging schedule of the AMoD vehicles and the proposed electricity prices); thus, the algorithm is privacy-preserving.

For ease of notation, we rewrite Equations (1)-(2) and (3)-(4) as, respectively,

$$\begin{aligned} f_{\text{TSO}}^{\text{eq}}(f_m, \lambda_m^{c, in}, \lambda_m^{t, c, out}, N_F) &= 0, \text{ (Eq. (1)-(2)), with dual vars. } \lambda_{\text{TSO}}^{\text{eq}}, \\ f_{\text{TSO}}^{\text{ineq}}(f_m, \lambda_m^{c, in}, \lambda_m^{t, c, out}, N_F) &\leq 0, \text{ (Eq. (3)-(4)), with dual vars. } \mu_{\text{TSO}}^{\text{ineq}}. \end{aligned}$$

Analogously, we rewrite Equation (6b) and Equations (6c)-(6f) as, respectively,

$$\begin{aligned} f_{\text{ISO}}^{\text{eq}}(f_m, \theta, p) &= 0, \text{ (Eq. (6b)), with dual variables } \lambda_{\text{ISO}}^{\text{eq}}, \\ f_{\text{ISO}}^{\text{ineq}}(f_m, \theta, p) &\leq 0, \text{ (Eq. (6c)-(6f)), with dual variables } \mu_{\text{ISO}}^{\text{ineq}}. \end{aligned}$$

We consider a partial Lagrangian relaxation of Problem (8), that is,

$$\begin{aligned} & \underset{f_m, \lambda_m^{c, in}, \lambda_m^{t, c, out}, N_F, \theta, p}{\text{minimize}} && V_T T_M(f_m) + V_D D_v(f_m) + V_B(f_m) + C_G(p) \\ & && + \lambda_{\text{ISO}}^{\text{eq}} f_{\text{ISO}}^{\text{eq}}(f_m, \theta, p) + \mu_{\text{ISO}}^{\text{ineq}} f_{\text{ISO}}^{\text{ineq}}(f_m, \theta, p), && (11a) \end{aligned}$$

$$\text{subject to} \quad f_{\text{TSO}}^{\text{eq}}(f_m, \lambda_m^{c, in}, \lambda_m^{t, c, out}, N_F) = 0, \quad (11b)$$

$$f_{\text{TSO}}^{\text{ineq}}(f_m) \leq 0. \quad (11c)$$

The TSO and the ISO iteratively optimize Problem (11) with respect to their own decision variables for a fixed value of the Lagrangian multipliers $\lambda_{\text{ISO}}^{\text{eq}}$ and $\mu_{\text{ISO}}^{\text{ineq}}$. Specifically, at step k of the iterative procedure, the TSO solves:

$$\begin{aligned} \underset{f_m^k, \lambda_m^{c,in,k}, \lambda_m^{t,c,out,k}, N_F^k}{\text{minimize}} \quad & V_T T_M(f_m^k) + V_D D_v(f_m^k) + V_B(f_m^k) \\ & + \lambda_{ISO}^{eq,k-1} f_{ISO}^{eq}(f_m^k) + \mu_{ISO}^{ineq,k-1} f_{ISO}^{ineq}(f_m^k), \end{aligned} \quad (12a)$$

$$\text{subject to} \quad f_{TSO}^{eq}(f_m^k, \lambda_m^{c,in,k}, \lambda_m^{t,c,out,k}, N_F^k) = 0, \quad (12b)$$

$$f_{TSO}^{ineq}(f_m^k) \leq 0. \quad (12c)$$

Minimizing the last two terms of Equation (12a) is equivalent to minimizing the cost of electricity V_E with prices $(\lambda_{ISO}^{eq,k-1} + \mu_{ISO}^{ineq,k-1})$. That is,

$$\arg \min_{f_m^k} \lambda_{ISO}^{eq,k-1} f_{ISO}^{eq}(f_m^k) + \mu_{ISO}^{ineq,k-1} f_{ISO}^{ineq}(f_m^k) = \arg \min_{f_m^k} V_E.$$

Thus, Problem (12) is equivalent to the Vehicle Routing and Charging Problem (5).

Analogously, at step k , the ISO solves

$$\underset{\theta^k, p^k}{\text{minimize}} \quad C_G(p^k) + \lambda_{ISO}^{eq,k-1} f_{ISO}^{eq}(\theta^k, p^k) + \mu_{ISO}^{ineq,k-1} f_{ISO}^{ineq}(\theta^k, p^k).$$

The Lagrangian multipliers are then updated by the ISO as

$$\begin{aligned} \lambda_{ISO}^{eq,k} &= \lambda_{ISO}^{eq,k-1} + \alpha_k \left(f_{ISO}^{eq}(f_m^k, \theta^k, p^k) \right), \\ \mu_{ISO}^{ineq,k} &= \max \left(0, \mu_{ISO}^{ineq,k-1} + \alpha_k \left(f_{ISO}^{ineq}(f_m^k, \theta^k, p^k) \right) \right), \end{aligned}$$

for an appropriately chosen step size α_k (see Lemma IV.2 below), and the TSO is informed of the new proposed price of electricity (i.e., the new value of the sum of the Lagrange multipliers).

Note that the ISO only needs to know the TSO's proposed charging schedule to compute $f_{ISO}^{eq}(f_m^k, \theta^k, p^k)$ and $f_{ISO}^{ineq}(f_m^k, \theta^k, p^k)$; in particular, the TSO does not need to disclose the customers' demand or the planned vehicle routes. Conversely, the ISO only needs to inform the TSO of the proposed price of electricity: the generation costs and the power demands remain private.

The next lemma proves that, if the step size α_k is "small enough," the algorithm converges.

Lemma IV.2 (Convergence of the dual decomposition algorithm). *If the step size α_k is chosen so that*

$$\begin{aligned} 0 < \alpha^k < 2 \left[- (V_T T_M(f_m^*) + V_D D_v(f_m^*) + V_B(f_m^*) + C_G(p^*)) \right. \\ & \left. + (V_T T_M(f_m^k) + V_D D_v(f_m^k) + V_B(f_m^k) + C_G(p^k)) \right] \\ & / \left(\|f_{ISO}^{eq}(f_m^k, \theta^k, p^k)\|^2 + \|f_{ISO}^{ineq}(f_m^k, \theta^k, p^k)\|^2 \right), \end{aligned} \quad (13)$$

then the dual decomposition algorithm converges to the optimal solution to Problem (8).

Proof. The proof follows immediately from Proposition 6.3.1 in [32]. \square

Note that the optimal value of Problem (8), $(V_T T_M(f_m^*) + V_D D_v(f_m^*) + V_B(f_m^*) + C_G(p^*))$, is not known. In practical applications, a small, fixed α^k and an appropriate stopping criterion should be used to ensure convergence.

V. NUMERICAL EXPERIMENTS

We study a hypothetical deployment of a P-AMoD system to satisfy medium-distance commuting needs in the Dallas-Fort Worth metroplex, with the primary objective of investigating the interaction between such a 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. The AMoD system is assumed to service 30% of all commuting trips, a scenario capturing low to medium penetration of AMoD. Census tracts in the metroplex are aggregated in 25 districts, as shown in Figure 2. We only consider trips starting and ending in different districts: 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 [33]. 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.

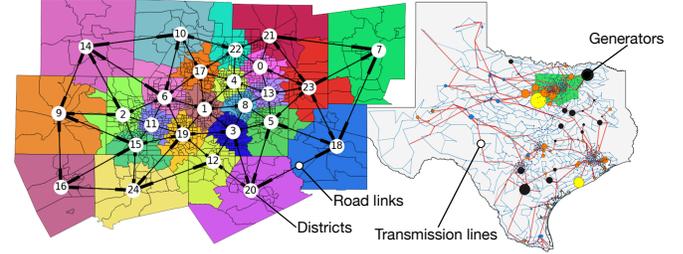


Fig. 2. Left: Census tracts and simplified road network for Dallas-Fort Worth. Right: Texas power network model (from [34]).

The battery capacity and power consumption of the EVs are modeled after the 2017 Chevrolet Bolt [35]. The cost of operation of the vehicles, excluding electricity costs, is \$0.16/mile (6.55¢/mile for maintenance and 9.46¢/mile for mileage-based depreciation), in accordance with AAA guidelines [36]. The fleet consists of 150,000 vehicles, i.e. 1 AMoD vehicle for every 2.67 customers, similar to the 2.6 ratio in [2]. 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 [34] 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 [37]. The model is also time-invariant; to model the time evolution of power demand 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.

For each study, we consider three different levels of battery depreciation. In the first case, the battery replacement cost

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

	Baseline	\$15,734 battery		\$1,573 battery		No depreciation	
		P-AMoD	Uncoord.	P-AMoD	Uncoord.	P-AMoD	Uncoord.
Avg. customer travel time [h]	-	1.0277	1.0277	1.0277	1.0277	1.0277	1.0277
Total energy demand [GWh]	517.498	520.543	520.543	520.543	520.544	520.590	520.966
Total electricity expenditure [k\$]	39,617.36	39,847.18	39,865.34	39,847.22	40,552.90	39,488.93	39,519.98
w.r.t. baseline [k\$]		+229.82	+247.98	+229.83	+935.54	-128.43	-97.38
Avg. price in DFW [\$/MW]	78.75	78.68	78.79	78.69	82.23	76.89	77.12
TSO electricity expenditure [k\$]	-	228.86	237.04	228.90	258.36	228.55	408.18

is \$15,734 (corresponding to the list price of a Chevrolet Bolt battery) and vehicles' batteries are fully depreciated over 1,000 charge-discharge cycles, in line with the performance of current battery technology. In the second case, the battery replacement cost is set to one tenth of the current one (or, equivalently, the vehicles' battery life is 10,000 cycles). In the third case, battery depreciation is neglected.

Table I and Figure 3 show the results. 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. Specifically, with current battery technology, coordination causes a 7.3% reduction in the TSO's electricity expenditure compared to the uncoordinated case, corresponding to savings of \$9M per year (assuming two commuting cycles per day and 250 work days per year). As battery prices are reduced ten-fold, the urgency of coordination between AMoD systems and the power network increases. In absence of coordination, the TSO's attempts to greedily charge and return power to the grid backfire, resulting in a four-fold increase in the TSO's electricity bill, a 4.4% increase in the unit price of electricity in the Dallas-Fort Worth area, and an additional expenditure of \$935k per day, or \$467M per year, in electricity costs borne by all power network customers. Conversely, coordination between the TSO and the ISO ensures that the unit price of electricity in the Dallas-Fort Worth area remains the *same* as in the baseline case, and results in savings of \$14.7M/year for the TSO compared to the uncoordinated case. A further reduction in the replacement cost of the batteries allows coordination between the AMoD system and the power network to reduce the *total* expenditure for electricity by \$128k 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. In the uncoordinated case, the presence of the TSO also reduces the overall electricity expenditure by \$97k/cycle compared to the baseline case - however, the reduction is offset by a \$180k/cycle increase in the TSO's own electricity bill compared to the coordinated case.

Who benefits from the reduction in energy expenditure? From the last two rows in Table I, one can see that, in the case where no depreciation is considered, the average price of electricity in the P-AMoD case is 2.37% lower than in the uncoordinated case in Dallas-Fort Worth (corresponding to savings of \$ 147M/year for Dallas-Fort Worth power network customers, excluding the TSO). The energy expenditure of the TSO in the P-AMoD case is 44% lower than in the uncoordinated case (a saving of \$180k per commuting cycle, corresponding to close to \$90M/year). Finally, electricity cus-

tomers outside of Dallas experience a small reduction of 0.23% 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. Figure 3 shows this phenomenon in detail for the

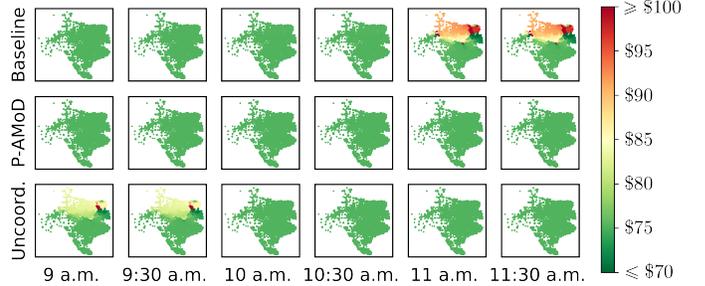


Fig. 3. LMPs in Texas between 9 a.m. and 11:30 a.m. The presence of the AMoD fleet can reduce locational marginal prices; coordination between the TSO and the ISO can yield a further reduction. A battery replacement cost of \$1,573 is considered.

scenario where the battery replacement cost is \$1,573. 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 result 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 Appendix A 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 minutes on commodity hardware and returns integral solutions that are directly amenable to control of P-AMoD systems. The algorithm allows us to perform agent-based simulations that provide further insights into the value of P-AMoD.

VI. CONCLUSIONS AND FUTURE WORK

In this paper we studied the interaction between an AMoD system and the electric 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 (Equation (10)), making realistic problems amenable to efficient numerical optimization on commodity hardware. We showed that the jointly optimal solution to the P-AMoD problem can be enforced as a general equilibrium, and we proposed a distributed privacy-preserving algorithm that allows agents to compute the market-clearing prices that enforce such an equilibrium without sharing private information about customer requests, generation costs, or power demands. We applied our model and algorithms to a case study of an AMoD deployment in Dallas-Fort Worth, TX. The case study showed that, depending on the maturity and cost of battery technology, 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; conversely, lack of coordination can result in large increases in power prices for power network customer and AMoD operators alike. These results are corroborated by agent-based simulations presented in the Appendix.

This work opens multiple avenues of research. First, 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. Second, we will extend the AMoD model to capture other modes of provision of service, including heterogeneous fleets where vehicles may differ in size, seating capacity, and battery capacity, and ride-pooling mechanisms where multiple customers with similar origins and destinations can travel in the same vehicle. 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] World Health Organization. (2014) 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.
- [2] K. Spieser, K. Treleven, 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.
- [3] R. Sioshansi, "OR Forum—modeling the impacts of electricity tariffs on plug-in hybrid electric vehicle charging, costs, and emissions," *Operations Research*, vol. 60, no. 3, pp. 506–516, 2012.
- [4] 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*, vol. 4, no. 4, pp. 863–875, 2017.
- [5] S. W. Hadley and A. A. Tsvetkova, "Potential impacts of plug-in hybrid electric vehicles on regional power generation," *The Electricity Journal*, vol. 22, no. 10, pp. 56–68, 2009.
- [6] R. Zhang and M. Pavone, "Control of robotic Mobility-on-Demand systems: A queuing-theoretical perspective," *Int. Journal of Robotics Research*, vol. 35, no. 1–3, pp. 186–203, 2016.
- [7] R. Iglesias, F. Rossi, R. Zhang, and M. Pavone, "A BCMP network approach to modeling and controlling autonomous mobility-on-demand systems," *Int. Journal of Robotics Research*, 2018, in Press.
- [8] M. Pavone, S. L. Smith, E. Frazzoli, and D. Rus, "Robotic load balancing for Mobility-on-Demand systems," *Int. Journal of Robotics Research*, vol. 31, no. 7, pp. 839–854, 2012.
- [9] F. Rossi, R. Zhang, Y. Hindy, and M. Pavone, "Routing autonomous vehicles in congested transportation networks: Structural properties and coordination algorithms," *Autonomous Robots*, vol. 42, no. 7, pp. 1427–1442, 2018.
- [10] R. Zhang, F. Rossi, and M. Pavone, "Model predictive control of Autonomous Mobility-on-Demand systems," in *Proc. IEEE Conf. on Robotics and Automation*, 2016.
- [11] J. Alonso-Mora, S. Samaranayake, A. Wallar, E. Frazzoli, and D. Rus, "On-demand high-capacity ride-sharing via dynamic trip-vehicle assignment," *Proceedings of the National Academy of Sciences*, vol. 114, no. 3, pp. 462–467, 2017.
- [12] M. W. Levin, K. M. Kockelman, S. D. Boyles, and T. Li, "A general framework for modeling shared autonomous vehicles with dynamic network-loading and dynamic ride-sharing application," *Computers, Environment and Urban Systems*, vol. 64, pp. 373 – 383, 2017.
- [13] D. Fiedler, M. Čertický, J. Alonso-Mora, and M. Čáp, "The impact of ridesharing in mobility-on-demand systems: Simulation case study in Prague," in *Proc. IEEE Int. Conf. on Intelligent Transportation Systems*, 2018, in Press.
- [14] N. Rotering and M. Ilic, "Optimal charge control of plug-in hybrid electric vehicles in deregulated electricity markets," *IEEE Transactions on Power Systems*, vol. 26, no. 3, pp. 1021–1029, 2011.
- [15] K. Turitsyn, N. Simitov, S. Backhaus, and M. Chertkov, "Robust broadcast-communication control of electric vehicle charging," in *IEEE Int. Conf. on Smart Grid Communications (SmartGridComm)*, 2010.
- [16] 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*, vol. 3, no. 4, pp. 1767–1778, 2012.
- [17] D. Goeke and M. Schneider, "Routing a mixed fleet of electric and conventional vehicles," *European Journal of Operational Research*, vol. 245, no. 1, pp. 81–99, 2015.
- [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*, vol. 26, no. 6, pp. 1325–1350, 2016.
- [19] L. Wang, A. Lin, and Y. Chen, "Potential impact of recharging plug-in hybrid electric vehicles on locational marginal prices," *Naval Research Logistics*, vol. 57, no. 8, pp. 686–700, 2010.
- [20] W. Kempton and J. Tomić, "Vehicle-to-grid power fundamentals: Calculating capacity and net revenue," *Journal of Power Sources*, vol. 144, no. 1, pp. 268–279, 2005.
- [21] M. E. Khodayar, L. Wu, and Z. Li, "Electric vehicle mobility in transmission-constrained hourly power generation scheduling," *IEEE Transactions on Smart Grid*, vol. 4, no. 2, pp. 779–788, 2013.
- [22] R. K. Ahuja, T. L. Magnanti, and J. B. Orlin, *Network Flows: Theory, Algorithms and Applications*. Prentice Hall, 1993.
- [23] J. G. Wardrop, "Some theoretical aspects of road traffic research," *Proc. of the Institution of Civil Engineers*, vol. 1, no. 3, pp. 325–362, 1952.
- [24] D. S. Kirschen and G. Strbac, *Fundamentals of Power System Economics*, 1st ed. John Wiley & Sons, 2004.
- [25] J. Glover, M. Sarma, and T. Overbye, *Power System Analysis and Design*, 5th ed. Cengage Learning, 2011.
- [26] W. W. Hogan, "Markets in real electric networks require reactive prices," in *Electricity Transmission Pricing and Technology*. Dordrecht: Springer Netherlands, 1996, ch. 7.
- [27] B. Stott, J. Jardim, and O. Alsaç, "DC power flow revisited," *IEEE Transactions on Power Systems*, vol. 24, no. 3, pp. 1290–1300, 2009.
- [28] R. P. O'Neill, T. Dautel, and E. Krall, "Recent ISO software enhancements and future software and modeling plans," Federal Energy Regulatory Commission, Tech. Rep., 2011.
- [29] 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.
- [30] H. D. Mittelmann. (2016) Decision tree for optimization software. <http://plato.asu.edu/guide>. [Online]. Available: <http://plato.asu.edu/guide>
- [31] G. Wang, M. Negrete-Pincetic, A. Kowli, E. Shafiepoorfar, S. Meyn, and U. V. Shanbhag, "Dynamic competitive equilibria in electricity markets," in *Control and optimization methods for electric smart grids*. Springer, 2012.
- [32] D. Bertsekas, *Nonlinear programming*, 2nd ed. Athena Scientific, 1999.
- [33] U.S. Dept. of Transportation, "Revised departmental guidance on valuation of travel time in economic analysis," Tech. Rep., 2015.
- [34] Illinois Center for a Smarter Electric Grid (ICSEG). (2016) Texas 2000-June 2016 synthetic power case. Information Trust Inst. Univ. of Illinois at Urbana-Champaign, Coordinated Science Laboratory. [Online]. Available: <http://icseg.iti.illinois.edu/synthetic-power-cases/texas2000-june2016/>
- [35] Chevrolet. (2017) Bolt EV. Available at <http://www.chevrolet.com/bolt-ev-electric-vehicle>. Retrieved on June 5, 2017.
- [36] AAA Association Communication, "Your driving costs," American Automobile Association, Tech. Rep., 2017.
- [37] EIA, "Levelized cost and levelized avoided cost of new generation resources in the annual energy outlook 2017," U.S. Energy Information Administration, Tech. Rep., 2017.
- [38] —, "Electric power annual 2016," U.S. Energy Information Administration, Tech. Rep., 2018.

APPENDIX

A. Agent-based simulations of P-AMoD

In this appendix we present agent-based simulations to further explore the impact of P-AMoD on the electric power network. First, by leveraging the structural insights from the network flow optimization problem, along with a few mild assumptions, we devise a computationally efficient control algorithm that solves the P-AMoD Problem (10) in a receding-horizon fashion. Due to space limitations, we only provide a high-level description of the algorithm: a detailed description is provided in the Extended Version of this paper¹.

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 $t_{i \rightarrow j}$ and a required charge of $c_{i \rightarrow j}$. 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 $t + t_{i \rightarrow j}$ with charge $c - c_{i \rightarrow j}$:

$$\lambda_m^{t,c,\text{out}} = \begin{cases} \lambda_m^{c+c_{v_m \rightarrow w_m}, \text{in}} & \text{if } t_m = t - t_{v_m \rightarrow w_m} \\ 0 & \text{otherwise} \end{cases} \quad (14)$$

for all $t \in \{1, \dots, T\}$, $c \in \{1, \dots, C\}$, $m \in \{1, \dots, M\}$. 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. In order to adapt the problem for real-time control of AMoD systems, one last difficulty must be overcome. The output of the problem is, in general, fractional: therefore it can not directly be used for control of individual vehicles. To overcome this, control actions are computed by *sampling* the first time step of the fractional optimal solution. We refer the reader to the Extended Version¹ for a thorough discussion.

We assess the performance of the receding-horizon P-AMoD controller with an agent-based simulation based on the same case study considered in Section V. The case study is modified in two ways: First, the generation costs are based on the *marginal* cost of generation (from EIA estimates [38, Table 8.4]), to reproduce the strategic behavior of generator operators participating in a real-time electricity market. Second, in order to assess the performance of the real-time controller under heavy load, *all* 1,257,916 commuting trips in Dallas-Fort Worth are serviced by an AMoD fleet of 450,000 vehicles.

The uncoordinated controller may cause the power network to become unstable, causing the Economic Dispatch problem (6) to become infeasible. To account for this, we introduce slack variables in the power network balance equations (6b). The slack variables capture the ISO's ability to disconnect loads to preserve the stability of the power network; the cost associated with the slack variables captures the economic loss borne by ISO users during a blackout (denoted as "Value of

Lost Load" in the literature) and is set to \$6,000/MWh in accordance with ERCOT estimates.

The receding-horizon problem is solved every 5 minutes with a 4-hour lookahead and a 15-minute time step. The performance of the algorithm is compared with a baseline case where no vehicles are present and an uncoordinated receding-horizon controller that optimizes the AMoD system's operations under the assumption that electricity prices stay constant.

Table II shows the results. In absence of coordination, the

TABLE II
REAL-TIME ALGORITHM SIMULATION RESULTS (10 HOURS).

	Baseline	P-AMoD	Uncoord.
Avg. cust. travel time [h]	-	1.585	1.591
Tot. energy demand [GWh]	500.00	507.71	508.02
Blackouts [MWh]	0	0	109.19
Tot. elec. expenditure, excl. TSO [k\$]	15,067	15,067	17,845
Avg. price in DFW [\$/MWh]	30.136	30.136	47.345
TSO tot. elec. expenditure [k\$]	-	232.61	4,446.61

AMoD system causes rolling blackouts in Dallas-Fort Worth: the Economic Dispatch problem is infeasible for 86 of the 600 minutes considered in the simulation, and overall 109.19 MWh of power are not delivered to end users. The average electricity price in Dallas-Fort Worth is \$47.35/MWh, 57% higher than in the baseline case; across Texas, the average price of electricity is \$43.88/MWh, and the total electricity expenditure for power network customers is almost 48% higher compared to the case where no vehicles are present. The TSO's expenditure is over 19 times higher compared to the coordinated case. Conversely, the P-AMoD system is able to ensure that the unit price of electricity (and therefore the expenditure of power network customers) in Dallas-Fort Worth and across Texas remains the *same* as in the case where no vehicles are present, despite the 4.82% increase in power demand in the Dallas-Fort Worth region. Thus, coordination between the AMoD system and the power network is vital to ensuring the stability of the power network. In absence of coordination, mass deployment of AMoD systems can heavily destabilize the power network, resulting in blackouts and excessive electricity prices; conversely, coordination is able to ensure that power prices remain constant despite the increase in power demand.

In contrast with the results in Section V, the AMoD system is not able to achieve a *reduction* in electricity prices. This is not entirely unexpected, as we use a comparatively short 4-hour lookahead which does not allow the system to fully exploit the daily variations in power demand. An important direction of future research is to explore how the lookahead time affects the tradeoff between computational complexity, economic savings, and robustness of the algorithm to inaccuracies in demand forecasting.

The receding-horizon P-AMoD problem was solved in an average of 66s and a maximum of 190s; thus, the algorithm is amenable to closed-loop control of large-scale systems.

B. Proofs of all theorems

Proof of Lemma III.2. The proof is constructive. First we leverage the flow decomposition algorithm to decompose the bundled customer flow in a collection of path flows; next, we assign each path flow to a customer request; finally, we merge the path flows assigned to each request to obtain a feasible customer flow. We assume without loss of generality that no two customer requests have the same origin node $v_m \in \mathcal{V}_R$, destination node $w_m \in \mathcal{V}_R$, and departure time $t_m \in \{1, \dots, T\}$. Since customer routes are approximated as

¹Available at <http://arxiv.org/abs/1709.04906v3>

a network flow, if two or more such requests exist, they can be equivalently represented by a single request with intensity equal to the sum of the original requests' intensities.

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 [22, Ch. 3.5] can decompose the bundled customer flow into path flows. Specifically, the algorithm computes a collection of path flows $\mathcal{P} = \{f_p(\mathbf{u}, \mathbf{v})\}_{p, (\mathbf{u}, \mathbf{v}) \in \mathcal{E}}$ such that, for every edge $(\mathbf{u}, \mathbf{v}) \in \mathcal{E}$, $\sum_p f_p(\mathbf{u}, \mathbf{v}) = f_{B, d_B}(\mathbf{u}, \mathbf{v})$. Each path flow $p \in \mathcal{P}$ has a single origin node $\mathbf{v} \in \mathcal{V}$ and destination node $\mathbf{w} \in \mathcal{V}$ with $v_{\mathbf{w}} = d_B$. Next, we assign each path flow to a customer request $(v_m, d_B, t_m, \lambda_m)$. Specifically, we decompose the path flows \mathcal{P} in a collection of disjoint sets $\{\mathcal{P}_m\}_m$ such that $\cup_{m=1}^M \mathcal{P}_m = \mathcal{P}$ and $\mathcal{P}_m \cap \mathcal{P}_{m'} = \emptyset$ for all $m, m' \in \{1, \dots, M\}$. To do so, we assign all the path flows whose origin node belongs to the set $\{\mathbf{v} = (v_m, t_m, c)\}_{c=1}^C$ to request m . By assumption, no two requests with the same destination d_B can have the same origin location v_m and departure time t_m : thus, every path flow is assigned to exactly one customer request m . The sum of the intensities of the path flows $p \in \mathcal{P}_m$ is λ_m ; this property follows immediately from Equations (9a) and (9b). Finally, the customer flow for customer request $(v_m, d_B, t_m, \lambda_m)$ is obtained as the sum of the path flows in \mathcal{P}_m . By construction, each path flow satisfies Equation (1a). Since the sum of the path flows equals λ_m , Equation (1b) is also satisfied by the sum of the path flows. This concludes the proof. \square

Proof of Theorem III.3. (i) *The bundled P-AMoD problem admits a feasible solution if the P-AMoD problem admits a feasible solution.* Consider a feasible solution to the P-AMoD problem. For each destination node, define the bundled flow as the sum of the customer flows for customers directed to that destination: $f_{B, d_B} = \sum_{m: w_m = d_B} f_m$ for all $d_B \in \mathcal{B}$.

It is easy to verify that the resulting network flow satisfies Equation (9). Also, the customer flows satisfy Equations (2), (3), (4), (6), and (7) and, for every edge, by construction $\sum_{d_B \in \mathcal{B}} f_{B, d_B} = \sum_{d_B \in \mathcal{B}} \sum_{m: w_m = d_B} f_m = \sum_m f_m$. Therefore the set of bundled customer flows $\{f_{B, d_B}\}_{d_B \in \mathcal{D}}$ satisfies Equations (2), (3), (4), (6), and (7) where each instance of $\sum_{m \in [1, M]} f_m$ is replaced by $\sum_{d_B \in \mathcal{D}} f_{B, d_B}$.

(ii) *The P-AMoD problem admits a feasible solution if the bundled P-AMoD problem admits a feasible solution.* Lemma III.2 shows that, if there exists a set of bundled flows that satisfy Problem (10), then there exists a set of customer flows that satisfy Equation (1). Furthermore, for each edge, $\sum_m f_m = \sum_{d_B \in \mathcal{B}} f_{B, d_B}$. Since the bundled flows satisfy the modified version of Equations (2), (3), (4), (6), and (7), the customer flows also satisfy them.

(iii) *The bundled P-AMoD problem and the P-AMoD problem have the same optimal value.* Due to Lemma III.2, $\sum_m f_m = \sum_{d_B \in \mathcal{B}} f_{B, d_B}$. The claim follows from the definition of the cost in Problem (10). \square

Proof of Theorem IV.1. The optimal solution to the P-AMoD problem also maximizes the revenue of the power generators if locational marginal pricing is used [31, Sec. 3]. Thus, we focus on showing that the optimal solution to the P-AMoD problem is also an optimal solution to the TSO's problem (5).

The KKT stationarity conditions for the P-AMoD Problem (8) for variables $\{f_m, \lambda_m^{c, \text{in}}, \lambda_m^{t, c, \text{out}}, N_F\}$ are:

$$\begin{aligned} & \frac{\partial(V_T T_M + V_D D_V + V_B)}{\partial f_m(\mathbf{v}, \mathbf{w})} + \lambda_{\text{TSO}}^{\text{eq}} \cdot \frac{\partial f_{\text{TSO}}^{\text{eq}}}{\partial f_m(\mathbf{v}, \mathbf{w})} + \mu_{\text{TSO}}^{\text{ineq}} \cdot \frac{\partial f_{\text{TSO}}^{\text{ineq}}}{\partial f_m(\mathbf{v}, \mathbf{w})} + \\ & \lambda_{\text{ISO}}^{\text{eq}} \cdot \frac{\partial f_{\text{ISO}}^{\text{eq}}}{\partial f_m(\mathbf{v}, \mathbf{w})} + \mu_{\text{ISO}}^{\text{ineq}} \cdot \frac{\partial f_{\text{ISO}}^{\text{ineq}}}{\partial f_m(\mathbf{v}, \mathbf{w})} = 0, \quad \forall m \in \{0, \dots, M\}, (\mathbf{v}, \mathbf{w}) \in \mathcal{E}, \end{aligned} \quad (15a)$$

$$\lambda_{\text{TSO}}^{\text{eq}} \cdot \frac{\partial f_{\text{TSO}}^{\text{eq}}}{\partial \lambda_m^{c, \text{in}}} = 0, \quad \forall c \in \{0, \dots, C\}, m \in \{0, \dots, M\}, \quad (15b)$$

$$\lambda_{\text{TSO}}^{\text{eq}} \cdot \frac{\partial f_{\text{TSO}}^{\text{eq}}}{\partial \lambda_m^{t, c, \text{out}}} = 0, \quad \forall c \in \{0, \dots, C\}, t \in \{1, \dots, T\}, m \in \{0, \dots, M\}, \quad (15c)$$

$$\lambda_{\text{TSO}}^{\text{eq}} \cdot \frac{\partial f_{\text{TSO}}^{\text{eq}}}{\partial N_F(\mathbf{v})} = 0, \quad \forall \mathbf{v} \in \mathcal{V}. \quad (15d)$$

For a given set of variables $\{\theta^*, p^*\}$, the KKT conditions for Problem (5) are

$$\begin{aligned} & \frac{\partial(V_T T_M + V_D D_V + V_B)}{\partial f_m(\mathbf{v}, \mathbf{w})} + \frac{\partial(V_E)}{\partial f_m(\mathbf{v}, \mathbf{w})} + \lambda_{\text{TSO}}^{\text{eq}} \cdot \frac{\partial f_{\text{TSO}}^{\text{eq}}}{\partial f_m(\mathbf{v}, \mathbf{w})} + \\ & \mu_{\text{TSO}}^{\text{ineq}} \cdot \frac{\partial f_{\text{TSO}}^{\text{ineq}}}{\partial f_m(\mathbf{v}, \mathbf{w})} = 0, \quad \forall m \in \{0, \dots, M\}, (\mathbf{v}, \mathbf{w}) \in \mathcal{E}, \end{aligned} \quad (16a)$$

$$\lambda_{\text{TSO}}^{\text{eq}} \cdot \frac{\partial f_{\text{TSO}}^{\text{eq}}}{\partial \lambda_m^{c, \text{in}}} = 0, \quad \forall c \in \{0, \dots, C\}, m \in \{0, \dots, M\}, \quad (16b)$$

$$\lambda_{\text{TSO}}^{\text{eq}} \cdot \frac{\partial f_{\text{TSO}}^{\text{eq}}}{\partial \lambda_m^{t, c, \text{out}}} = 0, \quad \forall c \in \{0, \dots, C\}, t \in \{1, \dots, T\}, m \in \{0, \dots, M\}, \quad (16c)$$

$$\lambda_{\text{TSO}}^{\text{eq}} \cdot \frac{\partial f_{\text{TSO}}^{\text{eq}}}{\partial N_F(\mathbf{v})} = 0, \quad \forall \mathbf{v} \in \mathcal{V}. \quad (16d)$$

The second term in Equation (16a) is

$$\frac{\partial(V_E)}{\partial f_m(\mathbf{v}, \mathbf{w})} = 1_{(\mathbf{v}, \mathbf{w}) \in \mathcal{E}_S} p(\mathbf{v}, \mathbf{w}) \delta c_{v_{\mathbf{v}}},$$

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

Leveraging Equation (7), the last two terms in Equation (15a) can be rewritten as

$$\begin{aligned} & \lambda_{\text{ISO}}^{\text{eq}} \frac{\partial f_{\text{ISO}}^{\text{eq}}}{\partial f_m(\mathbf{v}, \mathbf{w})} + \mu_{\text{ISO}}^{\text{ineq}} \frac{\partial f_{\text{ISO}}^{\text{ineq}}}{\partial f_m(\mathbf{v}, \mathbf{w})} = \sum_{l \in \mathcal{B}} \sum_{t=1}^T \left[\left(\lambda_{\text{ISO}}^{\text{eq}}(l, t) + \mu_{\text{ISO}}^{\text{ineq}}(l, t) \right) \cdot \right. \\ & \left. \left(1_{(\mathbf{v}, \mathbf{w}) \in M_{P, G}^+(l, t)} + 1_{(\mathbf{v}, \mathbf{w}) \in M_{P, G}^-(l, t)} \right) \right] J_C \delta c_{v_{\mathbf{v}}}. \end{aligned}$$

Every edge $(\mathbf{v}, \mathbf{w}) \in \mathcal{E}_S$ corresponds to a single load node $l \in \mathcal{B}$: $v_{\mathbf{v}} = \mathcal{M}_{\text{PR}}(l)$ at a single time $t = t_{\mathbf{v}}$. Thus, the expression above can be rewritten as

$$\lambda_{\text{ISO}}^{\text{eq}} \frac{\partial f_{\text{ISO}}^{\text{eq}}}{\partial f_m(\mathbf{v}, \mathbf{w})} + \mu_{\text{ISO}}^{\text{ineq}} \frac{\partial f_{\text{ISO}}^{\text{ineq}}}{\partial f_m(\mathbf{v}, \mathbf{w})} = J_C \delta c_{v_{\mathbf{v}}} \left(\lambda_{\text{ISO}}^{\text{eq}}(l_{v_{\mathbf{v}}}, t_{v_{\mathbf{v}}}) + \mu_{\text{ISO}}^{\text{ineq}}(l_{v_{\mathbf{v}}}, t_{v_{\mathbf{v}}}) \right),$$

where $l_{v_{\mathbf{v}}}$ is such that $v_{\mathbf{v}} = \mathcal{M}_{\text{PR}}(l_{v_{\mathbf{v}}})$.

The vector $(\lambda_{\text{ISO}}^{\text{eq}} + \mu_{\text{ISO}}^{\text{ineq}})$ denotes the locational marginal price of energy at each bus in the power network. That is,

$$p(\mathbf{v}, \mathbf{w}) = J_C \left(\lambda_{\text{ISO}}^{\text{eq}}(l_{v_{\mathbf{v}}}, t_{v_{\mathbf{v}}}) + \mu_{\text{ISO}}^{\text{ineq}}(l_{v_{\mathbf{v}}}, t_{v_{\mathbf{v}}}) \right),$$

where $l_{v_{\mathbf{v}}}$: $v_{\mathbf{v}} = \mathcal{M}_{\text{PR}}(l_{v_{\mathbf{v}}})$. (Note that $p(\mathbf{v}, \mathbf{w})$ is the price per discrete energy level, whereas $(\lambda_{\text{ISO}}^{\text{eq}}(l, t) + \mu_{\text{ISO}}^{\text{ineq}}(l, t))$ is the price per unit of energy). Therefore, Equations (16a) and (15a) are identical. As a result, the KKT conditions for the TSO's problem (5) are verified whenever the KKT conditions for the P-AMoD problem (8) are verified, and $\{f_m^*, \lambda_m^{c, \text{in}*}, \lambda_m^{t, c, \text{out}*}, N_F^*\}$ is an optimal solution to Problem (5) for fixed $\{\theta^*, p^*\}$.

In conclusion, $\{f_m^*, \lambda_m^{c, \text{in}*}, \lambda_m^{t, c, \text{out}*}, N_F^*\}$ is the solution to the TSO's Vehicle Routing and Charging Problem (5) if the prices are set according to LMPs. In addition, the generation schedule $\{p^*\}$ is the optimal (revenue-maximizing) schedule for self-interested power generators if the prices are set according to LMPs [31, Sec. 3]. That is, the set of variables $(\{f_m^*, \lambda_m^{c, \text{in}*}, \lambda_m^{t, c, \text{out}*}, N_F^*\}, \{\theta^*\}, \{p^*\})$ is a general equilibrium for the P-AMoD market. This concludes the proof. \square