Real-Time Distributed Control for Smart Electric Vehicle Chargers: From a Static to a Dynamic Study

Omid Ardakanian, S. Keshav, and Catherine Rosenberg
University of Waterloo
Technical Report CS-2013-20

***Abstract*—At high penetrations, uncontrolled electric vehicle (EV) charging has the potential to cause line and transformer congestion in the distribution network. Instead of upgrading components to higher nameplate ratings, we investigate the use of real-time control to limit EV load to the available capacity in the network. Inspired by rate control algorithms in computer networks such as TCP, we design a measurement-based, real-time, distributed, stable, efficient, and fair charging algorithm using the dual-decomposition approach. We show through extensive numerical simulations and power flow analysis on a test distribution network that this algorithm operates successfully in both static and dynamic settings, despite changes in home loads and the number of connected EVs. We find that our algorithm rapidly converges from large disturbances to a stable operating point. We show that in a test setting, for an acceptable level of overload, only 70 EVs could be fully charged without control, whereas up to around 700 EVs can be fully charged using our control algorithm. This compares well with the maximum supportable population of approximately 900 EVs. Our work also provides engineering guidelines for choosing the control parameters and setpoints in a distribution network.**

I. Introduction

At high penetration levels, uncontrolled electric vehicle charging can congest lines and transformers and cause voltage swings in the distribution system [1], [2]. Even at low penetration levels, uncontrolled charging can lead to congestion in certain neighbourhoods, due to a non-homogeneous distribution of EVs in the distribution network. Unrelieved congestion can overheat transformer windings and accelerate degradation of line and transformer insulation, leading to premature equipment failure. Although distribution system congestion can be relieved by upgrading system components piecemeal, this approach is both time-consuming and expensive. A more cost-effective alternative is for utility companies to directly control smart EV chargers[1] so that system components are rarely overloaded. This is the motivation for our work.

In prior work, EV chargers have been controlled using a schedule computed either the prior day (known as pre-dispatch scheduling) [1], [3]–[7] or in real-time [8]–[12]. Pre-dispatch scheduling approaches typically compute the charging schedule by solving a power flow problem. This requires precise estimates of non-EV loads, the points of connection of active EVs, their arrival and departure times, and the initial state of charge of their batteries. These parameters are difficult to predict accurately. Therefore, these approaches maintain a conservative operating margin to accommodate estimation uncertainties, which under-utilizes system resources.

In contrast, the real-time computation of charging schedules, which is the focus of this paper, achieves higher utilization by continuously adapting the charging rate of EV chargers to the *measured* available capacity of the network. In this approach, enabled by the widespread adoption of measurement and communication infrastructure in future distribution systems, line current and voltage measurements are sent from measurement nodes to control elements, and the control signals from these control elements are sent to EV chargers [13]. This allows EV chargers to use higher rates when there is available capacity, reducing these rates when the distribution network becomes congested. Note that during demand peaks, the available capacity of the network may not allow all EV chargers to charge at their maximum rate. Therefore, it is desirable to allocate the available capacity *fairly* among EV chargers. Thus, computing the set of EV charging rates can be viewed as an optimization problem whose solution is an allocation that simultaneously satisfies efficiency and fairness criteria.

Drawing on the design of congestion control protocols in packet-switched networks [14]–[17], our prior work [12] formulates a nonlinear convex optimization problem to obtain a charging rate allocation which is both proportionally fair [18], and scale-invariant Pareto optimal [19]. Solving the optimization problem allows us to obtain control rules that are implemented using dual decomposition [20] and the projected subgradient method [21]. In other words, we decompose the dual optimization problem into several subproblems, each solved independently and iteratively by an EV charger to adjust its charging rate. These subproblems are coordinated by a master problem through *congestion prices* [18], which are computed based on the congestion state of distribution lines and transformers and are periodically sent to EV chargers. We refer to this approach as *distributed control*.

Our prior work studied a *quasi-static* setting, where home loads were assumed to not change between snapshots. In this paper, we study our distributed control scheme in a *dynamic* setting, where home loads and the number of EVs being charged change over time. We validate using power flow analysis on a standard test distribution system that our control algorithm does not violate the operational limits of the distribution network. We investigate sensitivity of our control algorithm to arrivals and departures of EVs, the EV penetration

[1] A smart EV charger chooses a charging rate that not only optimizes battery life but also is responsive to control signals it receives.

level, the rated charge capacity of EV chargers, the choice of control parameters, and control setpoints. We make the following three contributions:

- We present a TCP-inspired measurement-based distributed control of EV charging and analyze its worst-case convergence.
- Using synthetic household load traces and an accurate power flow simulator, we numerically evaluate our control algorithm in a dynamic setting using extensive simulations, and study the performance sensitivity to the choice of control parameters.
- We provide engineering insights into the dynamic operation of the real-time distributed algorithm and discuss different design choices for control parameters to meet utility performance requirements.

This paper extends our prior work [12] in three ways. First, we validate the operation of the distributed control algorithm in a dynamic setting using power flow analysis, rather than in a quasi-static setting, as we had done earlier. Second, we analyze the convergence speed of the algorithm in the worst case. Finally, we provide several engineering insights.

II. Related Work

The potential impacts of introducing EVs into the distribution network have been explored extensively in the literature and many scheduling algorithms have been proposed to control EV charging load. Most existing work proposes centralized control of EV chargers. However, as discussed in a recent white paper [22], coordinating control at different levels becomes infeasible with such centralized control. This makes distributed control of EV charging and other responsive loads the better approach. Moreover, distributed control algorithms scale well with the size of the network and the number of EV chargers, and are robust to failure of a single measurement or control node. Therefore, we only review distributed control algorithms in the remainder of this section.

In recent work, Gan *et al.* [7] and Ma *et al.* [6] use distributed control to obtain a day-ahead charging schedule for EVs. Gan *et al.* formulate the EV charging control problem as an optimization problem with the objective of flattening the aggregate demand served by a transformer. A stochastic distributed control algorithm is proposed to find an approximate solution to this optimization problem. It is shown that this algorithm almost surely converges to one of the equilibrium charging profiles. In Reference [6], a decentralized algorithm is proposed to find the EV charging strategy that minimizes individual charging costs. It is shown that the optimal strategy obtained using this algorithm converges to the unique Nash equilibrium strategy when there is an infinite population of EVs. In the case of homogeneous EV populations, this Nash equilibrium strategy coincides with the valley-filling maximizing strategy (*i.e.*, the globally optimal strategy).

Our approach differs from these two approaches in two ways. First, their goal is to simply flatten the load served by the substation transformer, whereas we deal with line and transformer overloading in the entire distribution network. Second, these algorithms do not guarantee fair allocation of the available network capacity to EVs; this is an important property of our control mechanism.

Turning our attention to distributed real-time control, the closest line of work to ours is by Wen *et al.* [10], which selects a subset of connected EVs for charging at every timeslot so as to maximize user convenience subject to circuit-level demand constraints. This is formulated as a combinatorial optimization problem, and a centralized algorithm is proposed to solve a convex relaxation of this problem. Furthermore, a distributed algorithm is proposed to solve the optimization problem using the alternating direction method of multipliers for distributed optimization. Despite this similarity in using distributed optimization to schedule EV chargers in real-time, our approach differs in three ways. First, their focus is on satisfying user-specified charging requirements such as the charging deadline and the final state of charge, while our main goal is to balance efficiency and fairness[2]. Second, their formulation reflects only constraints imposed by two layers of a distribution network, whereas we model the entire distribution network, taking into account the capacities of all lines and transformers. Finally, unlike this work, they do not analyze the performance sensitivity of their control algorithm to the choice of control setpoints.

A recent paper by Fan [11] borrows the notion of congestion pricing from the Internet to reduce the peak load while providing weighted proportional fairness to end users. Exploiting the two way communications between the utility and users, congestion prices are sent to users, enabling them to adapt their demands to the capacity of the market in a fully distributed fashion. The user preference is modelled as a willingness to pay parameter, *i.e.*, the weight factor in the utility function of users. The proposed algorithm is then applied to EV charging to obtain a charging rate allocation. Interestingly, the total EV charging load varies with the range from which the weight factors can be chosen. Thus, the utility has to limit this range to ensure that the total load is not greater than the market capacity. The paper also studies the convergence behavior of the algorithm using both analysis and simulation results. However, unlike our approach, this paper does not model the distribution network and does not incorporate the capacity constraints of distribution lines and transformers, and the charging rate constraints of EV chargers.

In Reference [9], additive-increase-multiplicative-decrease (AIMD)-based charge control techniques are used for distributed control of EV charging. Interestingly, the authors study the problem from the user perspective rather than the utility perspective; they consider various scenarios and objective functions, and propose a separate AIMD-like algorithm for each scenario. Our approach differs from their approach in four ways. First, our congestion control algorithm deals with a network in which a set of EVs may share multiple lines and transformers, whereas their control algorithm assumes a single resource with a certain available capacity shared by all EVs. Second, they study several scenarios with different objective functions; however, we focus on implementation

[2]We believe that users have an incentive to lie about their deadlines to receive a higher charging rate if everyone pays the same.

and performance analysis of the control algorithm in a single scenario where the objective is to meet specific fairness and efficiency criteria. Third, we provide engineering guidelines to utilities for choosing control parameters and setpoints. Finally, our work is based on the theory of network utility maximization rather than an arbitrary choice of AIMD as the rate-control algorithm.

In [23] a control mechanism is designed to deal with the transformer overloading by modelling the transformer thermal limit as a constraint. Specifically, the authors formulate the EV charging problem as an open-loop centralized control problem with the objective of minimizing the SOC deviations from 1 and also minimizing the control effort subject to the capacity constraint of batteries and EV chargers, temperature constraint of the substation transformer, and the target SOC specified by EV owners. Using the dual decomposition method, an iterative price-coordinated implementation of this control mechanism is proposed which allows EV owners to compute their charging rate locally. To account for unexpected disturbances (i.e., fluctuations in the background demand), changing numbers of EVs, changing ambient temperature, and modelling errors they employ a receding-horizon feedback mechanism.

This work differs from ours in three ways. First, the proposed control method only deals with congestion at the substation transformer and does not address the overcapacity problem of feeders and other distribution transformers. Second, they do not consider fairness from the users perspective. Third, they do not model the distribution network, and do not use power flow analysis to study the operation of the algorithm in a test distribution network.

Note that none of the works discussed above evaluates the proposed solution using power-flow analysis, as we do.

This work builds on our prior work in this area. The idea of real-time distributed control of the EV charging was first introduced in [13] and this paper uses the control architecture discussed in [12]. Specific contributions of this work have already been clarified at the end of Section I.

III. Background and Assumptions

We study a radial distribution network that supplies both uncontrolled (home) and controlled (EV) loads, as illustrated in Figure 1. The topology of this network is encoded into a matrix R, where R_{sl} is 1 if charger s is downstream of line or transformer l, and is 0 otherwise.

A. Nameplate Rating and Setpoint

Every line or transformer in a distribution network has a *nameplate rating*. Equipment load must not exceed its nameplate rating over an extended period of time. We quantify distribution network congestion as the amount of energy transferred over and above the equipment's nameplate rating over a specified time period.

As part of our scheme, we allow a utility to associate a *setpoint* with every line or transformer. Our control goal is for the aggregate load (the sum of controlled and uncontrolled loads supplied by this equipment) to converge to this setpoint with only limited congestion, i.e., a limited number of excursions above the nameplate rating. Thus, a conservative utility can ensure a very low congestion level by choosing an appropriately low setpoint.

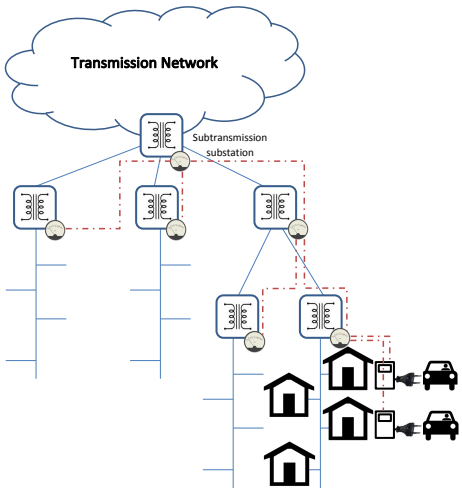

Fig. 1. An illustration of a smart distribution network consisting of MCC nodes, and communication links (dashed lines).

B. Measurement, Communication, and Control Nodes

We assume that lines and transformers are supplemented by *measurement, communication, and control* (MCC) nodes [13]. MCC nodes play three key roles in our proposed solution. First, they continuously measure the *congestion state* of their corresponding line or transformer, where the congestion state of a line or transformer is defined as the difference between its setpoint and its current loading level. Second, they compute *congestion prices* in each *signalling cycle* (see Section III-D). Finally, they send these congestion prices to EV chargers downstream to allow them to independently choose their rates (see Section VI).

As shown in Figure 1, the root MCC node is installed at the substation. It can send congestion signals on behalf of the (external) transmission and generation systems to reduce the EV charging load in response to generation shortfall or transmission network congestion. However, we do not consider these events in our work.

C. Assumptions

We now state the system assumptions that we make in developing our control algorithm.

- A1 The communication network is ubiquitous, broadband, reliable, and has a low latency.
- A2 MCC nodes can detect line or transformer overload sufficiently quickly that any transient overload is within system tolerances and the protection system is not invoked. This is true in nearly all distribution networks, where protection systems disconnect loads only when the overload is very large or persists for a long time.
- A3 It is not possible to infer congestion implicitly at EV chargers. Therefore, congestion must be explicitly signalled to them.
- A4 EVs only charge using EV chargers that are tamper-resistant and are under control of the electric utility. Thus,

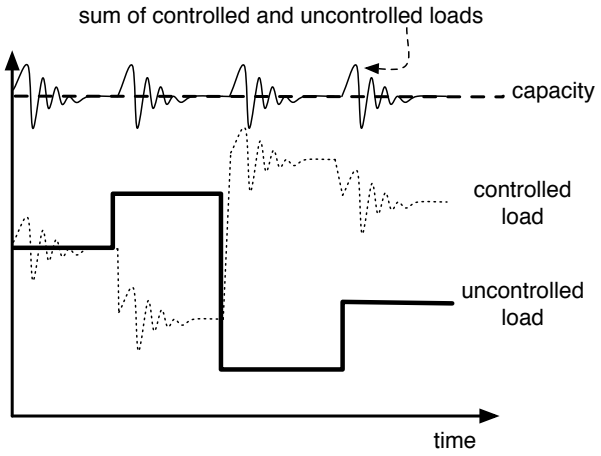

Fig. 2. The proposed control algorithm adapts charging rates of EV chargers to the available capacity of the network.

any control signal sent to them is assured of a cooperative response[3].

A5 An EV battery can be charged at any rate in the range $[0, m_s]$, where m_s is the maximum Amperage rating of its charger, independent of its state of charge[4].

A6 The power factor is close to unity and therefore reactive power flow can be neglected. This allows us to use a simple DC system model in our work. Moreover, distribution system losses are assumed to be negligible. We relax both assumptions in our simulation studies in Section IX.

These assumptions imply that it is feasible to design and implement a control algorithm that changes the EV charging rate rapidly in response to the congestion state of the distribution system.

D. System Operation

We briefly sketch the operation of our system. Every T_c milliseconds (the control timescale) the root MCC node initiates a signalling cycle by computing and sending its congestion price to its direct children. Upon receiving the congestion price(s), an intermediate MCC node computes and sends its own congestion price, using its latest recorded congestion state (as discussed in Section V-C1), along with the received price(s) to its children. Thus, EV chargers receive the set of congestion prices of all their parents. In Section V-C2, we explain how EV chargers use these congestion prices to choose their charging rate.

We now discuss the mathematical analysis used to obtain the congestion prices at each MCC node and the selection of charging rates at each EV charger.

IV. Optimization Problem

In this section we formulate the control problem as a centralized static constrained optimization problem. The global objective function is chosen such that the solution to this optimization problem also guarantees proportional fairness.

Our objective is to allocate the available capacity of the network fairly and efficiently among active EV chargers without overloading the distribution network. We adopt the notion of proportional fairness which is an axiomatically justified fairness criterion [19]. It can be shown that proportional fairness is achieved if we maximize the value of a global objective function defined as the sum of the logarithm of the utility function of the users [18]. Since the departure of EVs from homes and charging stations is non-deterministic, it is reasonable to assume that EV owners are greedy and prefer to finish charging their EVs as soon as possible. Hence, the utility of a user s is defined as the rate at which their EV is being charged, denoted x_s. For notational simplicity, we denote $\log(x_s)$ by $U_s(x_s)$. Observe that U_s is infinitely differentiable, increasing, and strictly concave on its domain.

Our optimization problem is therefore a maximization of the sum of the U_s of *active* chargers (*i.e.*, those chargers that are charging an EV), subject to physical constraints imposed by chargers, lines, and transformers. The constraint which corresponds to each line or transformer is that its loading level cannot exceed its setpoint. However, since each line or transformer supplies both homes and EV chargers[5], and the aggregate home load imposed on it along with the set of active EV chargers supplied by it change over time, we decompose the problem into a series of snapshots, where, in each snapshot, home loads are constant and a fixed number of EVs are plugged in to chargers. We then formulate an optimization problem for each snapshot and derive control rules from that problem.

Consider the n^{th} snapshot of the system in which the aggregate home load imposed on the line or transformer l is denoted by h_l^n, and the set of active chargers is denoted by $\mathcal{S}^n$. The optimization problem is:

$$\max_{x} \sum_{s\in\mathcal{S}^n} U_s(x_s) \tag{1}$$
$$\text{subject to} \quad 0 \le x_s \le m_s \quad \forall s \in \mathcal{S}^n$$
$$y_l \le \xi_l - h_l^n \quad \forall l \in \mathcal{L},$$

where ξ_l is the setpoint of l, y_l is the total EV charging load imposed on l (*i.e.*, $y_l = \sum_{s:R_{sl}=1} x_s$), and $\mathcal{L}$ is the set of distribution lines and transformers equipped with MCC nodes. This problem is a convex optimization problem because it maximizes an objective function that is the sum of concave functions (and is therefore concave), and each constraint defines a convex set. Note that we refer to the second inequality constraint of (1) as the *coupling constraint*.

In the next section we obtain the dual problem and apply the dual decomposition method to obtain a set of decoupled subproblems. We then design a distributed algorithm that solves these subproblems locally and independently.

V. Controller Design

The centralized optimization problem formulated in the previous section can be solved using a distributed approach.

[3]Note that our control algorithm would continue to operate if we assume cooperative response only from a subset of users; the uncooperative loads would act (and be considered) as uncontrolled home loads.

[4]With some battery technologies, the charging rate decreases as the state of charge increases. We do not consider this in our present analysis.

[5]For simplicity we ignore active power losses in the distribution network.

The distributed approach has three key advantages over the centralized approach. First, it gives autonomy to local controllers thereby increasing robustness of the control system. Second, it is more scalable. Third, it decreases the overall latency of control because control decisions are made locally.

Our plan is, therefore, to design a distributed control algorithm by solving the Lagrangian dual of the centralized optimization problem. We apply the dual decomposition method to obtain a set of decoupled subproblems that are controlled at the higher level by a master problem through congestion prices. The proposed algorithm requires solving the master problem and these subproblems in an iterative fashion. From a control theory standpoint, solutions to these problems constitute our controls and congestion prices are the feedback.

A. Dual Problem

Consider the Lagrangian relaxation of the optimization problem (1):

$$g(\lambda) = \max_{0 \preceq x \preceq m} \{\sum_{s \in \mathcal{S}} \log x_s + \sum_{l \in \mathcal{L}} \lambda_l (c_l^n - y_l)\}, \tag{2}$$

where $c_l^n = \xi_l - h_l^n$ denotes the available capacity of l in the n^{th} snapshot, $x = (x_1, \ldots, x_{|\mathcal{S}|})$ is a vector of charging rates, $m = (m_1, \ldots, m_{|\mathcal{S}|})$ is a vector of the charge capacity of EV chargers, $\lambda = (\lambda_1, \ldots, \lambda_{|\mathcal{L}|})$ is a vector of Lagrangian multipliers associated with the coupling constraints, and $\preceq$ is the vector inequality operator. Thus, the dual problem is

$$\begin{aligned} &\min_{\lambda} \max_{0 \preceq x \preceq m} \{\sum_{s \in \mathcal{S}} \log x_s + \sum_{l \in \mathcal{L}} \lambda_l (c_l^n - y_l)\} \\ &\text{subject to} \qquad \lambda_l \geq 0 \qquad \forall l \in \mathcal{L}, \end{aligned} \tag{3}$$

which is equivalent to

$$\begin{aligned} &\min_{\lambda} \left\{ \sum_{l \in \mathcal{L}} \lambda_l c_l^n + \max_{0 \preceq x \preceq m} \{\sum_{s \in \mathcal{S}} f_s(x_s; \lambda)\} \right\} \\ &\text{subject to} \qquad \lambda_l \geq 0 \qquad \forall l \in \mathcal{L}, \end{aligned} \tag{4}$$

where

$$f_s(x_s; \lambda) = \log x_s - x_s \sum_{l: R_{sl}=1} \lambda_l, \tag{5}$$

Note that we do not introduce dual variables for the constraints $0 \leq x_s \leq m_s$; hence, the maximization over x is restricted to $0 \leq x_s \leq m_s$ for all values of s (see Section 3.4.2 in [21]). In the above equation, $f(x; \lambda)$ represents f as a function of x parameterized by λ. Since $f_s(x_s; \lambda)$ is the sum of two concave functions of x_s, it is also concave and has a unique maximum.

We note that (4) is derived from (3) by using the following equation.

$$\sum_{l \in \mathcal{L}} \left(\lambda_l \sum_{s: R_{sl}=1} x_s \right) = \sum_{s \in \mathcal{S}} \left(x_s \sum_{l: R_{sl}=1} \lambda_l \right) = x R \lambda^T$$

Importantly, in our formulation, strong duality holds because all inequality constraints are affine. Therefore, we can write the following KKT optimality conditions

$$\hat{x}_s = \left[\frac{1}{\sum_{l: R_{sl}=1} \hat{\lambda}_l} \right]_0^{m_s} \qquad \forall s \in \mathcal{S} \tag{6}$$

$$\hat{\lambda}_l (\hat{y}_l - c_l^n) = 0 \qquad \forall l \in \mathcal{L} \tag{7}$$

$$\hat{y}_l \leq c_l^n \qquad \forall l \in \mathcal{L} \tag{8}$$

$$0 \leq \hat{x}_s \leq m_s \qquad \forall s \in \mathcal{S} \tag{9}$$

$$\hat{\lambda}_l \geq 0 \qquad \forall l \in \mathcal{L} \tag{10}$$

where $\hat{x}$ and $\hat{\lambda}$ are the unique optimizers of the Lagrangian dual problem. The first condition says that the gradient of Lagrangian vanishes at the optimal point, and the second condition, *i.e.*, the complementary slackness condition, implies that either the optimal Lagrangian multiplier is zero, or the corresponding line or transformer is *fully utilized*, *i.e.*, the line or transformer loading reached its nominal setpoint. Combining the first three conditions gives us the following relation between $\hat{x}$ and $\hat{\lambda}$.

$$\hat{x}_s = min \left\{ \frac{1}{\sum_{l: R_{sl}=1} \hat{\lambda}_l}, m_s \right\} \tag{11}$$

B. Dual Decomposition

Writing the Lagrangian dual problem in the form of (4) reveals its hidden decomposition structure [20]. Specifically, each EV charger can locally solve a subproblem given by

$$\max_{0 \leq x_s \leq m_s} f_s(x_s; \lambda), \tag{12}$$

provided that it knows the sum of the Lagrangian multipliers corresponding to the lines and transformers that are supplying its load. It turns out that Lagrangian multipliers play the role of congestion prices (or *shadow prices* [14]) in our problem.

These subproblems are controlled by a master problem by means of congestion prices. The master problem is responsible for updating the congestion prices and can be written in the following form

$$\min_{\lambda \succeq 0} \left\{ \sum_{l \in \mathcal{L}} \lambda_l c_l^n + \sum_{s \in \mathcal{S}} f_s(\hat{x}_s; \lambda) \right\}. \tag{13}$$

where $f_s(\hat{x}_s; \lambda)$ is the optimal value of (12). Observe that the objective function of the master problem is linear in λ and its derivative with respect to a Lagrangian multiplier is given by

$$\frac{\partial g}{\partial \lambda_l}(\lambda) = c_l^n - y_l$$

C. Control Laws

Our approach is to solve the dual optimization problem using a distributed algorithm which has two separate parts. The first part adjusts congestion prices of lines and transformers by periodically measuring the available capacity and solving the master problem at each MCC using the gradient projection method. The second part updates the charging rates of EVs by solving the subproblems.

In the following we derive control laws for updating congestion prices and adjusting charging rates by solving the master

problem and the subproblems respectively. These control laws constitute the distributed algorithm outlined in Section VI. In Section VII, we specify sufficient conditions for convergence of this algorithm to primal and dual optimal values.

1) A Control Law for Updating the Congestion Price: Since the dual function is differentiable, we can adopt the gradient method with a projection onto the positive orthant to solve the master problem (13). The following algorithm updates congestion prices in each iteration in opposite direction to the gradient of the dual function.

$$\lambda_l(t+1) = \max\{\lambda_l(t) - \kappa(c_l^n - y_l(t)), 0\} \quad \forall l \in \mathcal{L} \quad (14)$$

Here κ is a sufficiently small positive constant which determines the responsiveness and stability of control. Note that it is not necessary to estimate c_l^n and y_l separately at an MCC node to compute $c_l^n - y_l$. This is because $c_l^n - y_l$ is equal to the congestion state of l, *i.e.*, the total line or transformer loading subtracted from its setpoint, and the congestion state is what being measured by the corresponding MCC node.

2) A Control Law for Adjusting the Charging Rate: We denote the latest congestion price vector that an EV charger has received by $\lambda(\bar{t})$. The subproblem (12) can be easily solved by finding the stationary point of $f_s(x_s; \lambda)$.

$$f_s'(x_s(t); \lambda(\bar{t})) = \frac{1}{x_s(t)} - \sum_{l:R_{sl}=1} \lambda_l(\bar{t}) \overset{\text{set to}}{=} 0 \rightarrow$$

$$x_s(t) = \min\left\{\frac{1}{\sum_{l:R_{sl}=1}\lambda_l(\bar{t})}, m_s\right\} \quad (15)$$

Note that $x_s(t)$ is the rate of EV charger s for the interval $[t, t+1)$, and adjusting the charging rates impacts the loading of upstream feeders and transformers immediately[6]. More specifically, $y_l(t)$ is given by

$$y_l(t) = \sum_{s:R_{sl}=1} x_s(t) \qquad \forall l \in \mathcal{L} \quad (16)$$

Note that the unit of time in (14) and (15) is T_c milliseconds, and therefore $\bar{t}$ equals $t - \frac{d}{T_c}$ because congestion prices are received by EV chargers after d milliseconds, which is an upper bound on the one-way latency from an MCC node and its downstream EV chargers.

VI. Distributed Charging Control Algorithm

We now describe the algorithms that operate at the MCC nodes and at EV chargers and implement the control laws derived in Section V.

Our distributed charging control algorithm measures the congestion state of a line or a transformer and computes the corresponding congestion price based on (14). This price is sent to descendant EV chargers every T_c milliseconds (see Algorithm 1).

After receiving congestion prices from upstream MCC nodes, every charger computes its charging rate using (15) and starts charging at this rate (see Algorithm 2).

[6]There is a fundamental difference between congestion control protocols in the Internet and our EV charge control protocol. In computer networks, when traffic sources change their rates it is only reflected on link utilization after a delay, known as the forward delay. However, there is no forward delay in our problem as power flows in the grid at the speed of light.

Algorithm 1: Congestion price update at MCC node l with setpoint ξ_l

input: ξ_l, $\kappa(> 0)$

while *true* **do**
- Measure load
- congestion state $\leftarrow \xi_l -$ load
- price $\leftarrow \max\{$price $-\ \kappa \times$ congestion state$, 0\}$
- Send price along with all received prices to children
- Wait until the next **message from parent** or the next **clock tick** (if root)

end

Algorithm 2: Rate adjustment at EV charger s

input: m_s, new congestion prices

while *true* **do**
- $\lambda \leftarrow$ new congestion prices
- aggregate price $\leftarrow \sum_{l \in ascendants} \lambda_l$
- rate $\leftarrow \min\{\frac{1}{\text{aggregate price}}, m_s\}$
- Start charging the battery at rate
- Wait until the next **message from parent**

end

VII. Convergence Analysis

This section investigates the impact of control parameters (the gradient step size and the timescale of control) on the stability and the convergence speed of the algorithm. We first study the conditions under which the proposed distributed control algorithm converges to the solution of the centralized optimization problem (1) in a static setting, *i.e.*, no EVs arrive or depart and the change in the magnitude of uncontrollable loads is negligible, then study the worst-case rate of convergence.

A. Proof of Stability

Note that the primal optimum is equal to the dual optimum as strong duality holds. Therefore, in this setting, we only need to show that the distributed control algorithm converges to the solution of (3). We then verify convergence in a dynamic setting both by studying the worst-case change in home loads, and through extensive numerical simulations.

Let $\overline{L} := \max_s \sum_l R_{sl}$ be the length of the longest path from the substation to an EV charger, $\overline{S} := \max_l \sum_s R_{sl}$ be the maximum number of active EV chargers sharing a link, $\overline{m} := \max_s m_s$ be the maximum charging rate supported by EV chargers, and d be the maximum communication delay between the root MCC node and any EV charger.

Theorem 1. *Starting from any initial charging rates $0 \preceq x \preceq m$ and congestion prices $\lambda \succeq 0$, the distributed control algorithm converges to the primal-dual optimal values if*

(1) $T_c \geq d$
(2) $0 < \kappa < \kappa^* = \frac{2}{\overline{m}^2 \overline{LS}}$

Proof sketch: The first condition guarantees that the control action at each MCC node, i.e., a change in congestion price, is taken only after all the EV chargers have reacted to the

previous control action. In this case, the continuous time system reduces to the discrete-time system studied in [15] and our theorem reduces to Theorem 1 proved in that work. The second condition maps directly to the necessary condition for Theorem 1 in [15]. □

B. Convergence Speed in the Worst Case

We now prove that the control algorithm exhibits monotone convergence even in the worst case and exploit this behaviour to compute an upper bound on its convergence time.

Clearly, the control algorithm converges most slowly when there is the greatest need for a decrease in EV charging rates, coupled with the least possible decrease in the charging rate in each iteration. This is because each iteration of the control algorithm reduces the EV load by a multiplicative factor, where the factor corresponds to the sum of the non-negative congestion prices sent by upstream MCC nodes. Thus, the worst case is when all home loads are initially zero and all EVs charge at their maximum rate; subsequently, a *single* line or transformer is overloaded due to the increase in the aggregate home load to its maximum possible value. For simplicity, we assume that prior to the worst-case change in home load, control has stabilized, with the gradient step size equal to κ^*, and all EV chargers have the same charging capability, *i.e.*, $m_s = m$ for all s.

We first prove that system convergence after a worst-case change in home load is monotone.

1) Monotonic Convergence: Suppose that the worst-case change in home loads happens during the n^{th} snapshot that starts at t_0. Because the system is underloaded prior to the change, all congestion prices are zero at time t_0. Denote the line or transformer that overloads as l. Then, $y_l(t_0^+) > c_l^n$. We now prove that y_l converges *monotonically* to its setpoint.

The total EV charging load supplied by l is

$$\begin{aligned} y_l(t) &= \sum_{s\in S(l)} x_s(t) = \sum_{s:R_{sl}=1} \min\{\frac{1}{\sum_{l':R_{sl'}=1}\lambda_{l'}(t)}, m\} \\ &= \sum_{s:R_{sl}=1} \min\{\frac{1}{\lambda_l(t)}, m\} \\ &= \sum_{s:R_{sl}=1} \frac{1}{\lambda_l(t)} \qquad \text{if } \lambda_l(t) > \frac{1}{m} \\ &= \frac{|S(l)|}{\lambda_l(t)} \end{aligned} \tag{17}$$

where $|S(l)|$ is the number of EV chargers downstream of l. Note that the second line of (17) is derived from the first line since l is the only congested line or transformer, and therefore the congestion price of other lines and transformers is zero.

Observe that y_l remains constant for a few iterations after t_0 until λ_l exceeds $\frac{1}{m}$, starting from zero. Let t_s be the first time in this snapshot that the condition $\lambda_l(t) > \frac{1}{m}$ holds, then y_l starts decreasing when this price is received by EV chargers. The following equation can be derived from (17) for $t \geq t_s$:

$$\begin{aligned} y_l(t) - y_l(t+1) &= |S(l)|(\frac{1}{\lambda_l(t)} - \frac{1}{\lambda_l(t+1)}) \\ &= |S(l)|\frac{\lambda_l(t+1) - \lambda_l(t)}{\lambda_l(t)\lambda_l(t+1)} \end{aligned} \tag{18}$$

The following theorem states that monotonic convergence of our control is guaranteed in the worst case.

Theorem 2. *If, in a distribution network, the length of the longest path from the substation to an EV charger is at least 2, for all $t > t_0$, $y_l(t) \geq c_l^n$ if $y_l(t_0) \geq c_l^n$.*

Proof. We prove this theorem by contradiction. Suppose for some $\tilde{t} \geq t_s$ we have $y_l(\tilde{t}) \geq c_l^n$, and $y_l(\tilde{t}+1) < c_l^n$, or equivalently $y_l(\tilde{t}) - y_l(\tilde{t}+1) > y_l(\tilde{t}) - c_l^n$. From (18) we have:

$$|S(l)|\frac{\kappa(y_l(\tilde{t}) - c_l^n)}{\lambda_l(\tilde{t})(\lambda_l(\tilde{t}) + \kappa(y_l(\tilde{t}) - c_l^n))} > y_l(\tilde{t}) - c_l^n$$

The following inequality is obtained by canceling out the $y_l(\tilde{t}) - c_l^n$ term from both sides of the above inequality

$$\frac{|S(l)|}{\lambda_l(\tilde{t})} - \frac{\lambda_l(\tilde{t})}{\kappa} > y_l(\tilde{t}) - c_l^n$$

But $\frac{|S(l)|}{\lambda_l(\tilde{t})} - \frac{\lambda_l(\tilde{t})}{\kappa} \leq 0$ since

$$\begin{aligned} |S(l)|\kappa &= |S(l)|\frac{2}{m^2\overline{LS}} \\ &\leq \frac{2}{m^2\overline{L}} \\ &\leq \frac{1}{m^2} \qquad \text{since } \overline{L} \geq 2 \\ &\leq \lambda_l(\tilde{t})^2 \end{aligned}$$

Thus $y_l(\tilde{t}) - c_l^n < \frac{|S(l)|}{\lambda_l(\tilde{t})} - \frac{\lambda_l(\tilde{t})}{\kappa} \leq 0$. This contradicts our assumption that $y_l(\tilde{t}) - c_l^n \geq 0$. Therefore we conclude that $\tilde{t}$ does not exist, and y_l does not go below c_l^n. □

2) An Upper Bound: The monotonic convergence of the control algorithm enables us to compute an upper bound on the time that it takes until the algorithm converges to the region specified by $\pm\tau$ around the setpoint when l is overloaded at t_0 and $y_l(t_0) - c_l^n = \Delta$.

We first compute an upper bound on the length of the interval between t_0 and t_s, denoted δ_s, during which y_l does not change. We can write

$$\begin{aligned} \delta_s &\leq \left(\lfloor\frac{\frac{1}{m} - 0}{\kappa(y_l(t) - c_l^n)}\rfloor + 1\right) \times T_c \\ &= \left(\lfloor\frac{1}{m\kappa\Delta}\rfloor + 1\right) \times T_c \end{aligned} \tag{19}$$

Now let t_{conv} be the first time that the loading of l goes below the level $c_l^n + \tau$ and δ_{conv} be an upper bound on the length of the interval between t_{conv} and t_s. Since $y_l(t) - c_l^n \geq \tau$ for $t < t_{\text{conv}}$, we can write $\lambda_l(t) = \frac{|S(l)|}{y_l(t)} < \frac{|S(l)|}{c_l^n + \tau}$ for $t < t_{\text{conv}}$. From (18) we obtain:

$$y_l(t) - y_l(t+1) > \kappa\tau\frac{(c_l^n + \tau)^2}{|S(l)|} \tag{20}$$

for $t < t_{\text{conv}} - 1$.

To compute δ_{conv} we approximate the decaying decrease rate of the line or transformer loading (after t_s) with a constant

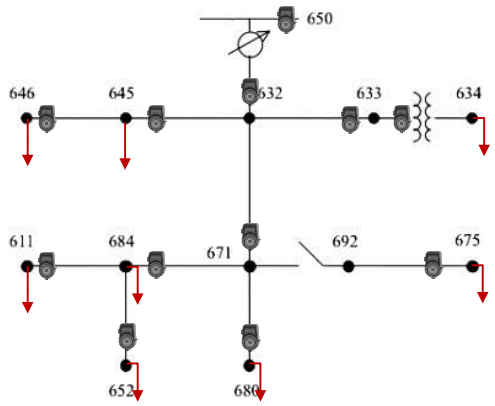

Fig. 3. The one-line diagram of our test distribution network. MCC nodes are shown as meters placed on the lines.

value equal to $y_l(t_{\text{conv}}) - y_l(t_{\text{conv}} - 1)$:

$$\delta_{\text{conv}} = \left(\lceil \frac{y_l(t_0) - (c_l^n + \tau)}{y_l(t_{\text{conv}} - 1) - y_l(t_{\text{conv}})} \rceil + 1 \right) \times T_c$$
$$= \left(\lceil \frac{(\Delta - \tau)|S(l)|}{\kappa \tau (c_l^n + \tau)^2} \rceil + 1 \right) \times T_c \quad (21)$$

Thus, $\delta_{\text{conv}} + \delta_s$ is an upper bound on the convergence time when l is overloaded at t_0.

VIII. Test Distribution System

We evaluate our control algorithm by means of power flow analysis using the Open Distribution System Simulator (OpenDSS) [24] on the standard IEEE 13-bus test feeder [25], a 4.16kV three-phase radial distribution system (Figure 3). In this section, we discuss the details of this distribution network as well as our approach to model home and EV loads.

This distribution network is supplied by a three-phase 5MVA transformer that steps down the transmission line voltage from 115kV to 4.16kV. For sake of simplicity, we treat buses as load aggregation points with directly connected home loads and EV chargers and do not model the transformers and feeders radiating from them, although that analysis would be a straightforward adaptation to what we describe below. Using the parameters provided in [25], we set the nameplate rating of every feeder to its ampacity at $50°C$. The nameplate rating of the substation transformer and the in-line transformer are also set to 5MVA and 500kVA respectively. We assume that all loads, including homes and EV chargers, are single-phase and are connected between a phase and neutral of load buses 634, 645, 646, 675, 680, 684, 652, 611. We further assume that EV chargers are identical, either Level 1 (a maximum load of 1.8kW) or Level 2 (a maximum load of 7.2kW), consume real power only, and the reactive power consumption of every home is 30% of its real power consumption, a conservative assumption.

A. MCC Nodes

MCC nodes measure the load at a phase conductor on a load bus (Figure 3). Similarly, the loading of the substation transformer and the in-line transformer, connected between buses 633 and 634, are also measured by an MCC node. All MCC nodes are interconnected using a communication network which forms a logical tree overlaid on the radial distribution network. Thus, for example the MCC node installed at phase B of bus 632 is the parent of MCC nodes installed at phase B of buses 633, 645, and 671, and the MCC node installed at the substation is the parent of MCC nodes installed at phases A, B, and C of bus 632.

B. Power Flow

To run a power flow study, we specify the amount of active and reactive power injected at every load bus. The OpenDSS simulator solves the optimization problem and provides the outputs power flows in different branches. Assuming that measurements from MCC nodes match the result of power flow calculations, we use these results in our control algorithm to obtain congestion prices. We compute the charging rate of EV chargers using the new prices sent by upstream MCC nodes, and update the aggregate EV charging load at every bus accordingly. We also update the aggregate home load at every bus using the synthetic load model described in Section VIII-C. This allows us to run a power flow study for the next iteration.

C. Home Load Model

We assume that our test distribution network supplies 3300 households, connected to load buses as described in Table I, and a finite population of EVs.

To study the fast timescale dynamics of our algorithm, we need fine-grained measurements of the household loads, which we lack. Therefore, we generate synthetic load traces using the Markov models developed in [26] for household electricity consumption during on-peak, mid-peak, and off-peak periods. These models are derived from fine-grained measurements of electricity consumption in 20 homes over four months. We then compute the aggregate home load imposed on the transformer by adding these loads, ignoring losses. Our simulations span over three days in winter and the corresponding aggregate home load is illustrated in Figure 7. The peak of the aggregate home load is 4.44MW and hence the distribution system is never congested over these three days in the absence of EVs.

D. EV Model

We assume that each household has at most one EV and a single EV charger that can charge only one EV at a time. We assume that the capacity of an EV battery is 24kWh (the capacity of a Nissan Leaf EV). We also assume that all EVs leave the system every day after 6am following a Poisson distribution with parameter μ_d, and return to the system after 4pm following a Poisson distribution with parameter μ_a, with fully discharged batteries. Thus, the number of charging EVs changes with time. Since the EV population is finite, a higher value of μ_a (or μ_d) creates a larger burst of arrivals around 4pm (or departures around 6am).

TABLE I

THE TOTAL HOME LOAD AND THE NUMBER OF EV CHARGERS CONNECTED TO A SPECIFIC PHASE OF A LOAD BUS IN THE STATIC SETTING, AND THE NUMBER OF HOMES AND THE PERCENTAGE OF EV POPULATION CONNECTED TO A SPECIFIC PHASE OF A LOAD BUS IN THE DYNAMIC SETTING.

	Bus	680			634			675			645		646		684		652	611
	Phase	a	b	c	a	b	c	a	b	c	b	c	b	c	a	c	a	c
Static	Agg. home load (kW)	600	600	600	100	100	100	400	400	400	100	100	300	300	100	100	200	200
	Num of chargers	80	80	80	40	40	40	40	40	40	40	40	40	40	40	40	40	40
Dynamic	Num of homes	450	450	450	50	50	50	300	300	300	50	50	200	200	50	50	150	150
	Percentage of chargers	10%	10%	10%	5%	5%	5%	5%	5%	5%	5%	5%	5%	5%	5%	5%	5%	5%

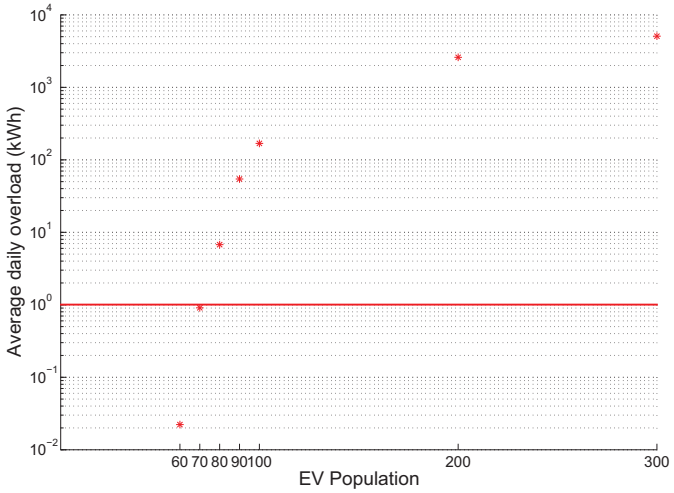

Fig. 4. Average daily overload versus EV population when EVs are charged by AC Level 2 chargers without control. Note that the Y-axis is logarithmic scale.

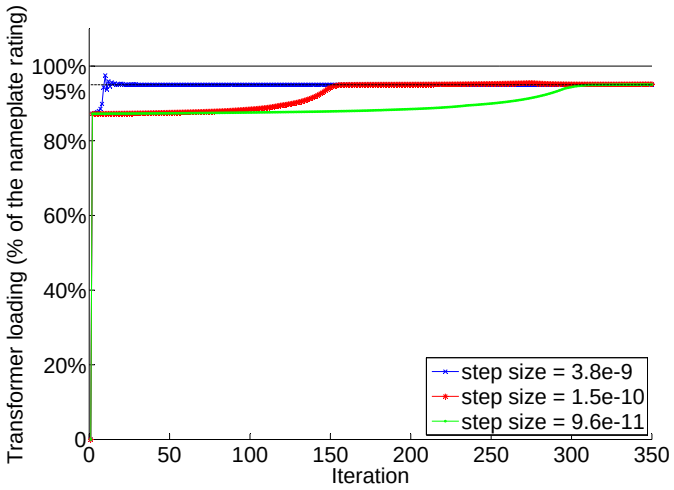

Fig. 5. The value of the step size κ determines how the loading of the substation transformer changes over time. Here the equipment setpoint is set to 95% of its nameplate rating.

IX. PERFORMANCE EVALUATION

This section uses power flow analysis to study:

- the effect of uncontrolled EV loads on a distribution network (Section IX-A),
- how long it takes for the distributed algorithm to converge to the chosen setpoint (Section IX-B),
- the dynamic behaviour of our system (Section IX-C), and
- the efficiency of our control system, compared to the best possible efficiency (Section IX-D).

We measure the *efficiency* by the average energy stored in an EV battery in a day. We measure the *congestion* of a line or a transformer by the amount of energy it supplied when it is above its nameplate rating level (we refer to this as the *overload* below).

A. The Need for Control

Figure 4 shows the effect of uncontrolled EV loads on the distribution network using Level 2 chargers. We find that a population of merely 90 EVs in a neighbourhood comprised of 3300 homes leads to a non-negligible overload of 54kWh/day. To avoid congestion without controlled charging, either the EV penetration level must be kept low or lines and transformers must be upgraded. For example, if the utility requires an overload of less than 1kWh/day, depicted by a horizontal line in Figure 4, without controlled charging, the EV population must be kept below 70 or 2.1%. Even when the overall penetration level is low, this may not be true for certain neighbourhoods.

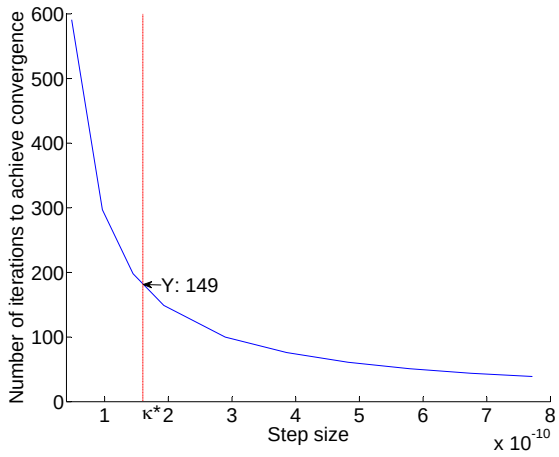

Fig. 6. The number of iterations to achieve convergence.

B. Rate of Convergence

We now study the number of iterations that it takes to achieve convergence for different values of κ, assuming that the setpoint of each line or transformer is set to be 95% of its nameplate rating, that EV chargers charge at Level 1, and that the home loads do not change. Table I summarizes our simulation scenario in this *static* setting.

In this scenario, the maximum charging rate is 1800 W, so the maximum step size for which the convergence of the algorithm is guaranteed is $\kappa^* = \frac{2}{1800^2 \times 800 \times 5} = 1.54 \times 10^{-10}$ (from Theorem 1). As we increase the value of κ, the control system transitions from an over-damped system to an under-damped system and eventually to an unstable system for

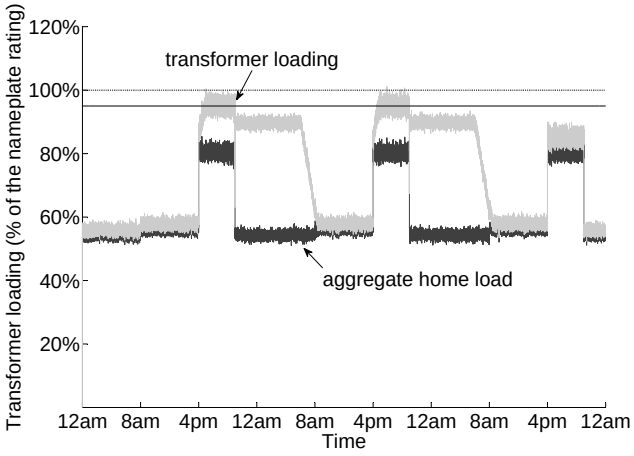

Fig. 7. Operation of the algorithm in a dynamic setting with 900 Level 1 EV chargers where the substation transformer's setpoint is 4.75MVA (the solid line) and the control timescale is 1 second. The aggregate home load is shown in dark grey and the overall load measured at the substation transformer in light grey.

$\kappa > 3.8 \times 10^{-9}$ which is larger than κ^*. Figure 5 shows how the loading of the substation transformer varies over different iterations for three different values of κ that are smaller than, equal to, and larger than κ^*.

The value of κ also controls the number of iterations it takes to achieve convergence, that is, when the loading of a line or transformer is within $\pm 1\%$ of its setpoint. Figure 6 shows that the number of iterations required for convergence deceases exponentially as we increase the value of κ. When the step size is equal to κ^*, it takes 149 iterations to achieve convergence (about 15 seconds if the control timescale is chosen to be 0.1s).

C. System Dynamics

We now investigate the scenario when the home loads and the number of EVs are time-varying, as discussed in Section VIII. In these experiments, we fix the value of the step size to be κ^*, and study six different values for the setpoint of the substation transformer: 4.75MVA, 4.8MVA, 4.85MVA, 4.9MVA, 4.95MVA, and 5MVA, two different control timescales: 1 second, and 0.5 second, and two possible charging levels: AC Level 1 (1.8 kW maximum) and AC Level 2 (7.2 kW maximum) [27]. We repeat each simulation 10 times, using 10 different arrival and departure times generated by setting $\mu_d = \mu_a = 0.1$ per second[7].

Figure 7 shows a single simulation trace of the overall load over time when the transformer setpoint is 4.75MVA, the control timescale is 1 second, and the EV population is 900. The home loads have evening peaks, with loads at night and at mid-day being roughly equal. Note that even when EVs are not charging, the transformer load is slightly higher than the aggregate home load due to line losses. When EVs are present, the overall load is close to the setpoint, with rare excursions above the nameplate rating. Each such excursion contributes to the *overload*. Clearly, the lower the setpoint, the lower this overload. To demonstrate this, Figure 8 shows the average daily overload versus the transformer's setpoint when the EV population is 500. It can be readily seen that the overload decreases with lower setpoints, but increases as we use slower control timescales. It also increases significantly with the charging level, and grows exponentially with the setpoint for a specific control timescale and a specific charging level.

[7]Thus, on average, one EV arrives and departs every 10 seconds after 4pm and 6am respectively. These arrival and departure patterns are intentionally chosen to stress test the control algorithm.

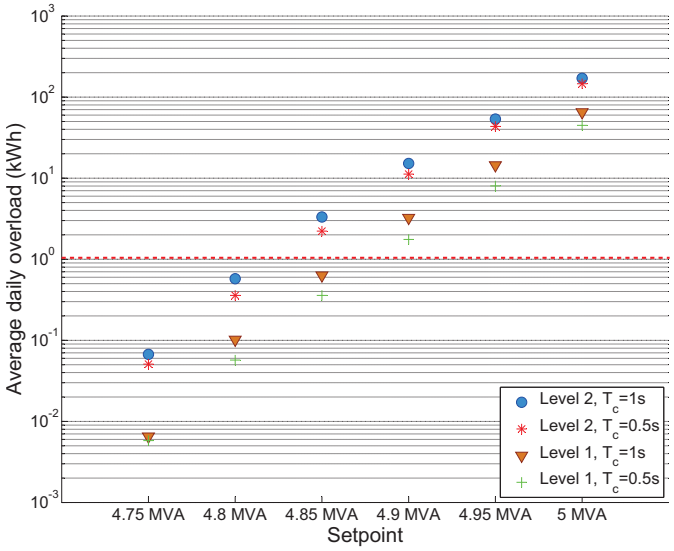

Fig. 8. Average daily overload versus setpoint when the EV population is 500. Note that the Y-axis is logarithmic scale.

D. Efficiency

We now study the efficiency of our control scheme. Consider the situation in which all EVs arrive at 4pm and stay in the system until 6am of the next day. We numerically compute an upper bound on the number of EVs that can be fully charged in this time interval by simply dividing the integral of the difference between the nameplate rating and the aggregate home load (after incorporating losses) over this interval by the battery capacity (*i.e.*, 24kWh). We find that a maximum of 900 EVs can be fully charged between 4pm and 6am of the next day under these ideal conditions. With our control algorithm, using Level 2 charging, if we set the setpoint of the substation transformer to 4.8MVA to obtain a very small overload (approximately 1kWh/day), up to around 700 EVs can be fully charged which compares well with the maximum of 900 EVs especially when recalling that without control we cannot charge more than 70 EVs to obtain the same level of overload.

X. Engineering Insights

This section provides guidelines for choosing the control parameters and setpoints based on the results of our simulations, assuming that the utility limits the amount of risk that is willing to take. For simplicity, we confine our study to the substation transformer which is the potential bottleneck in this network, although our study applies to arbitrary multi-level distribution networks. Thus, we only have one setpoint to select.

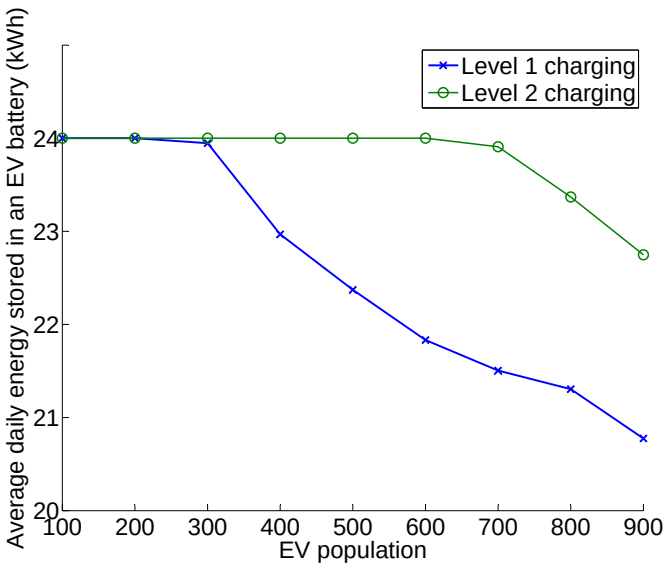

Fig. 9. Average energy stored in an EV battery in a day versus the EV population for AC Level 1 and Level 2 chargers when the control timescale is 1 second.

A. Choosing Control Knobs: κ and T_c

Recall that the gradient step size, κ, should be set to κ^* to improve responsiveness of control and ensure stability at the same time.

Choosing T_c is more complex. Faster control timescales reduce the overload but increase the communication overhead. If communication is not a bottleneck, then T_c should be set as small as possible. Of course, this is lower bounded by the communication delay, as stated by Theorem 1.

Communication overhead can be reduced by using multicast to send packets from an MCC node to its children. With multicasting we can implement the control algorithm by sending at most $\mathcal{O} = |\mathcal{S}| + |\mathcal{L}| - 1$ packets in every control interval, where $|\mathcal{S}|$ is the number of EV chargers and $|\mathcal{L}|$ is the number of MCC nodes. Hence, $\mathcal{O}$, $2 \times \mathcal{O}$, and $10 \times \mathcal{O}$ packets must be transmitted per second when the control timescale is 1 second, 0.5 second, and 0.1 second respectively. Thus the control timescale can be chosen properly given the communication medium and protocol.

B. Choosing Setpoints

Recall that transformer overloading increases the risk of equipment failure and outages. The choice of setpoints of lines and transformers depends on the amount of risk the utility is willing to take. Specifically, the utility chooses the equipment setpoint for a given EV population and a charging level such that the overload does not exceed an acceptable level.

For example, in the above scenario (where we have only one setpoint to select) the setpoint of the transformer could be as high as 4.85MVA to obtain an overload of 1kWh/day (depicted by a dashed line in Figure 8) when the EV population is 500 and all EV chargers are Level 1.

Figure 9 shows the average daily energy stored in an EV battery for different EV population sizes using Level 1 and Level 2 chargers when the setpoint is chosen to obtain an overload of less than 1kWh/day, and the control timescale is set to 1 second. Observe that EV batteries are fully charged using both Level 1 and Level 2 chargers when the EV population is less than 300 (Level 2 charging only reduces the charging time). However, when the EV population exceeds 300, the system becomes overly congested; in this regime, Level 2 charging is more beneficial to EV owners than Level 1 charging as it increases the efficiency. For example, when the EV population is 500 and all chargers are Level 2, the setpoint could be as high as 4.8MVA; this corresponds to on average 24kWh of energy transferred to an EV per day, which means that all batteries can be fully charged. Similarly, when the EV population is 500 and all chargers are Level 1, the setpoint could be as high as 4.85MVA; this corresponds to on average 22.37kWh of energy transferred to an EV per day.

XI. Conclusion

Our work addresses line and transformer congestion arising from uncontrolled charging of electric vehicles. Motivated by rate control in computer networks, we propose a real-time, distributed, stable, efficient, and fair charging algorithm based on the dual-decomposition approach. This algorithm scales well with the size of the network and the number of EVs. We show through extensive numerical simulations as well as power flow analysis on a test distribution network that this algorithm operates successfully in both static and dynamic settings, despite changes in home loads and the number of connected EVs. We analyze the sensitivity of this algorithm to the EV penetration level, the rated charge capacity of EV chargers, the choice of control parameters, and control setpoints. Based on this analysis, we provide guidelines for choosing these parameters in a distribution network.

Our work suffers from some limitations. The chief limitation is that it requires the installation of MCC nodes at all potential points of congestion in the distribution network. A second problem is that our control algorithm does not scale particularly well: the choice of κ, which controls system responsiveness, is upper-bounded by the number of EVs and the maximum EV charging rate. As these increase, κ and system responsiveness decrease. What is needed is a less conservative bound for κ that has better scaling properties. Third, our EV battery model is simple and ignores the relationship between the charging rate and the state-of-charge of the battery. Finally, our approach does not deal with other consequences of EV charging including large voltage swings and phase imbalance. We plan to address these limitations in future work.

References

[1] J. Lopes, F. Soares, and P. Almeida, "Integration of electric vehicles in the electric power system," *Proceedings of the IEEE*, vol. 99, no. 1, pp. 168–183, 2011.

[2] S. Shao, M. Pipattanasomporn, and S. Rahman, "Challenges of PHEV penetration to the residential distribution network," in *Power Energy Society General Meeting, IEEE*, 2009, pp. 1–8.

[3] C. Ahn, C.-T. Li, and H. Peng, "Optimal decentralized charging control algorithm for electrified vehicles connected to smart grid," *J. Power Sources*, vol. 196, no. 2, pp. 10 369 – 10 379, 2011.

[4] E. Sortomme, M. Hindi, S. MacPherson, and S. Venkata, "Coordinated charging of plug-in hybrid electric vehicles to minimize distribution system losses," *Smart Grid, IEEE Transactions on*, vol. 2, no. 1, pp. 198–205, 2011.

[5] K. Clement-Nyns, E. Haesen, and J. Driesen, "The impact of charging plug-in hybrid electric vehicles on a residential distribution grid," *Power Systems, IEEE Transactions on*, vol. 25, no. 1, pp. 371–380, 2010.

[6] Z. Ma, D. Callaway, and I. Hiskens, "Decentralized charging control of large populations of plug-in electric vehicles," *Control Systems Technology, IEEE Transactions on*, vol. 21, no. 1, pp. 67–78, 2013.

[7] L. Gan, U. Topcu, and S. Low, "Optimal decentralized protocol for electric vehicle charging," *Power Systems, IEEE Transactions on*, vol. 28, no. 2, pp. 940–951, 2013.

[8] S. Deilami, A. Masoum, P. Moses, and M. A. S. Masoum, "Real-time coordination of plug-in electric vehicle charging in smart grids to minimize power losses and improve voltage profile," *Smart Grid, IEEE Transactions on*, vol. 2, no. 3, pp. 456–467, 2011.

[9] S. Studli, E. Crisostomi, R. Middleton, and R. Shorten, "Aimd-like algorithms for charging electric and plug-in hybrid vehicles," in *International Electric Vehicle Conference (IEVC)*. IEEE, 2012, pp. 1–8.

[10] C.-K. Wen, J.-C. Chen, J.-H. Teng, and P. Ting, "Decentralized plug-in electric vehicle charging selection algorithm in power systems," *Smart Grid, IEEE Transactions on*, vol. 3, no. 4, pp. 1779–1789, 2012.

[11] Z. Fan, "A distributed demand response algorithm and its application to phev charging in smart grids," *Smart Grid, IEEE Transactions on*, vol. 3, no. 3, pp. 1280–1290, 2012.

[12] O. Ardakanian, C. Rosenberg, and S. Keshav, "Distributed control of electric vehicle charging," in *Proceedings of e-Energy'13*. ACM, 2013, pp. 101–112.

[13] ——, "Realtime distributed congestion control for electrical vehicle charging," *SIGMETRICS Perform. Eval. Rev.*, vol. 40, no. 3, pp. 38–42, Jan. 2012.

[14] F. P. Kelly, A. K. Maulloo, and D. K. H. Tan, "Rate control for communication networks: Shadow prices, proportional fairness and stability," *Journal of the Operational Research Society*, vol. 49, no. 3, pp. 237–252, 1998.

[15] S. H. Low and D. E. Lapsley, "Optimization flow control. I. Basic algorithm and convergence," *Networking, IEEE/ACM Transactions on*, vol. 7, no. 6, pp. 861–874, 1999.

[16] R. Srikant, *The Mathematics of Internet Congestion Control (Systems and Control: Foundations and Applications)*. Birkhauser, 2004.

[17] F. Paganini, Z. Wang, J. Doyle, and S. Low, "Congestion control for high performance, stability, and fairness in general networks," *Networking, IEEE/ACM Transactions on*, vol. 13, no. 1, pp. 43–56, 2005.

[18] F. Kelly, "Charging and rate control for elastic traffic," *European transactions on Telecommunications*, vol. 8, no. 1, pp. 33–37, 1997.

[19] H. Yaïche, R. R. Mazumdar, and C. Rosenberg, "A game theoretic framework for bandwidth allocation and pricing in broadband networks," *IEEE/ACM Trans. Networking*, vol. 8, no. 5, pp. 667–678, 2000.

[20] D. Palomar and M. Chiang, "A tutorial on decomposition methods for network utility maximization," *Selected Areas in Communications, IEEE Journal on*, vol. 24, no. 8, pp. 1439–1451, 2006.

[21] D. P. Bertsekas and J. N. Tsitsiklis, *Parallel and distributed computation: numerical methods*. Upper Saddle River, NJ, USA: Prentice-Hall, Inc., 1989.

[22] J. Taft and P. De Martini, "Cisco Systems – Ultra Large-Scale Power System Control Architecture," http://www.cisco.com/web/strategy/docs/energy/control_architecture.pdf.

[23] R. Hermans, M. Almassalkhi, and I. Hiskens, "Incentive-based coordinated charging control of plug-in electric vehicles at the distribution-transformer level," in *American Control Conference (ACC), 2012*, 2012, pp. 264–269.

[24] EPRI, "Simulation Tool OpenDSS," http://www.smartgrid.epri.com/SimulationTool.aspx.

[25] W. Kersting, "Radial distribution test feeders," in *Power Engineering Society Winter Meeting, 2001. IEEE*, vol. 2, 2001, pp. 908–912 vol.2.

[26] O. Ardakanian, S. Keshav, and C. Rosenberg, "Markovian models for home electricity consumption," in *Proc. ACM SIGCOMM Green Networking Workshop*, 2011.

[27] "SAE J1772 Standard," http://www.sae.org/smartgrid/chargingspeeds.pdf.