

# Distributed Control of Electric Vehicle Charging

Omid Ardakanian  
University of Waterloo  
oardakan@uwaterloo.ca

Catherine Rosenberg  
University of Waterloo  
cath@uwaterloo.ca

S. Keshav  
University of Waterloo  
keshav@uwaterloo.ca

## ABSTRACT

Electric vehicles (EVs) are expected to soon become widespread in the distribution network. The large magnitude of EV charging load and unpredictable mobility of EVs make them a challenge for the distribution network. Leveraging fast-timescale measurements and low-latency broadband communications enabled by the smart grid, we propose a distributed control algorithm that adapts the charging rate of EVs to the available capacity of the network ensuring that network resources are used efficiently and each EV charger receives a fair share of these resources. We obtain sufficient conditions for stability of this control algorithm in a static network, and demonstrate through simulation in a test distribution network that our algorithm quickly converges to the optimal rate allocation.

## Categories and Subject Descriptors

I.2.8 [Artificial Intelligence]: Problem Solving, Control Methods, and Search—*control theory*

## Keywords

Electric vehicle charging, congestion control, distributed control

## 1. INTRODUCTION

Unlike gasoline-powered cars, battery electric vehicles and plug-in hybrid electric vehicles, both referred to as ‘EVs’, are powered by electricity stored in their on-board battery which is charged when EVs are plugged in to chargers located at homes or public charging stations. EV chargers can impose a significant load on the distribution network: with AC Level 2 charging, EVs can be charged at up to 80A at 240V [2], a load of 19.2kW, whereas a typical North American home has an average load of only 1kW. Therefore, a single EV being charged at the peak Level 2 rate could impose an instantaneous load as large as that imposed by nearly

twenty average homes. Consequently, the large-scale introduction of EVs is likely to greatly affect the electrical grid’s distribution system [5, 7, 13]. It can cause overloading of distribution branches and transformers, and voltage drop at distant buses. Persistent overloading can cause damage to conductors, overheat transformers, and degrade their insulation. Excessive voltage drop can cause damage to electrical appliances. Both may eventually lead to a loss of reliability due to the invocation of the protection system.

The branch and transformer congestion problem is known as a major barrier to large-scale adoption of EVs [13]. Using lower-level charging does reduce the impact on the grid but only at the expense of greatly increasing the duration of the charging process and the inefficiency of the network due to under-utilization. A compromise solution is therefore to exploit elasticity of EV charging load to control charging rates such that bottleneck lines and transformers are fully utilized, and the line voltage level remains within a predetermined range.

Existing approaches to control EV charging load either use a central controller to coordinate charging [10, 13, 20] or cast the control problem in the form of a distributed optimization [3, 8, 15]; in both approaches the charging schedule is computed well ahead of time and decisions are made on the basis of system-level considerations, such as mitigating distribution system losses or maximizing the load factor. In the centralized approach, the central controller uses power flow analysis to compute a charging schedule that does not congest any part of the distribution network. This analysis requires an accurate model of the distribution network and the expected locations of all EVs. In many cases, such a model is either not available or not up-to-date. Critically, both approaches also need to predict the future demand from non-EV loads, the number of charging EVs at each time slot, their locations, and their initial state of charge. The safety margin built in to hedge against prediction errors makes both approaches quite conservative.

Inspired by the design of the Internet, which offers best-effort services to elastic applications that back off in case of congestion [19], our approach is to quickly adapt EV charging rates to the condition of the network [4]. Specifically, we propose a distributed control algorithm so that every charger can independently set its charging rate based on *congestion signals* it receives from measurement nodes installed on its path to the subtransmission substation. This algorithm ensures that EV chargers receive a *proportionally fair* [11] share of the available capacity of the distribution network, and lines and transformers are not overloaded. Therefore, an

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

*e-Energy*’13, May 21–24, 2013, Berkeley, California, USA.  
Copyright 2013 ACM 978-1-4503-2052-8/13/05 ...\$15.00.

EV can be charged at the maximum rate when its feeding branches and transformers are lightly loaded (*e.g.*, during off-peak hours), and it is charged at a relatively low rate when they are heavily loaded (*e.g.*, during on-peak hours). This is a best-effort service; hence, in the rare event that the grid is overly congested some EVs might not be fully charged by their deadlines. Note that this is in line with today’s practice where operators shed the load in favor of protecting their transmission and distribution assets.

We make the following three specific contributions:

- We formulate the EV charging control problem as an optimization problem where the objective function is chosen such that proportional fairness is achieved at its solution.
- We decompose the optimization problem into a set of distributively solvable subproblems and derive stable control laws by iteratively solving these subproblems.
- We illustrate using numerical simulations that this control converges in a small number of iterations to the solution of the utility maximization problem.

In this paper, we focus on a *static* network scenario in which the residential load is constant and a fixed number of EVs are connected to chargers. We explain in Section 3.2 that the dynamic case can be decomposed into a series of static snapshots. Thus, the control algorithm developed for a static scenario can also be used in a network with variable home loads and number of plugged-in EVs.

We make the critical assumption that the congestion level of every branch is measured in real time and is communicated to downstream EV chargers with a reasonably low delay. This requires an infrastructure which allows fast measurement and communication. The future smart grid is likely to have a considerable number of measurement and control devices that are interconnected by a ubiquitous low-latency broadband communication network [6]. This allows us to use fast time scale measurements and communication to rapidly adapt the charging rates of EVs to the available capacity of the network such that sustained overloading of branches is avoided [4].

## 2. RELATED WORK

Potential impacts of introducing a large number of EVs to the distribution network have been explored extensively in the literature and many scheduling algorithms have been proposed to shift the EV charging load to off-peak hours, thereby avoiding branch congestion and voltage drop in the distribution network. Most existing work suggests a centralized control for EV charging load. However, as discussed in a recent white paper [22], coordinating control at different levels becomes infeasible with such centralized control. This highlights the need for distributed control of EV charging and other responsive loads.

The closest lines of work to ours are by Gan *et al.* [8] and Ma *et al.* [15] which use distributed control to obtain a day-ahead charging schedule for EVs. In [8], it is assumed that the distribution transformers and EV chargers are instrumented with computation and communication devices, EVs are charged at a fixed rate, and the charging process cannot be interrupted. Based on these assumptions, the EV charging control problem is formulated as a discrete 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 [15], 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 the approach of these two papers in three ways. First, we control charging of EVs in real time, whereas they compute a day-ahead charging schedule based on predictions. Second, their goal is to simply flatten the load at a single point in the network, whereas we deal with line and transformer overloading in the entire distribution network. Third, these algorithms do not guarantee fair allocation of available network capacity to EVs while this is an important property of our control mechanism.

The idea of real time distributed control of the EV charging was first introduced in a vision paper [4]; a measurement and signalling architecture was proposed, and three possible distributed congestion control schemes for EV charging were outlined. This paper builds upon the architecture proposed in [4]. We make more precise assumptions about the underlying system, and propose a distributed algorithm that would enable fair, timely, and efficient charging of EVs starting from a static centralized optimization problem.

Our EV charging control problem has the same mathematical formulation as the rate control problem which was extensively studied in the context of the Internet [12, 14, 17, 21, 23, 24] and was used to analyze stability and other properties of the TCP congestion control protocol (see [4] for a comprehensive comparison of congestion control in a packet-switched network and the power distribution network). The common goal is to determine the available resources and allocate these resources among users to maximize a global objective function of the users’ utilities that takes fairness into account. Control rules are found by solving the optimization problem in a distributed fashion exploiting the hidden decomposition structure of the optimization problem (see [18] for an introduction to network utility optimization and decomposition theory). Depending on the choice of the objective function, the solution to this maximization problem provides different notions of fairness, namely proportional fairness, max-min fairness, minimum potential delay fairness, and the more general notion of utility proportional fairness.

## 3. BACKGROUND AND ASSUMPTIONS

The electrical grid consists of generation, transmission, and distribution systems. The electricity generated by power plants is transmitted over long distances by the transmission network, *i.e.*, a mesh network of high voltage lines and step up transformers. Near demand centers, the voltage is stepped down to the primary distribution voltage at sub-transmission substations. The distribution system, which is the focus of this paper, is responsible for delivering electricity from these substations to consumers. A radial distribution system has a tree structure and is comprised of nodes

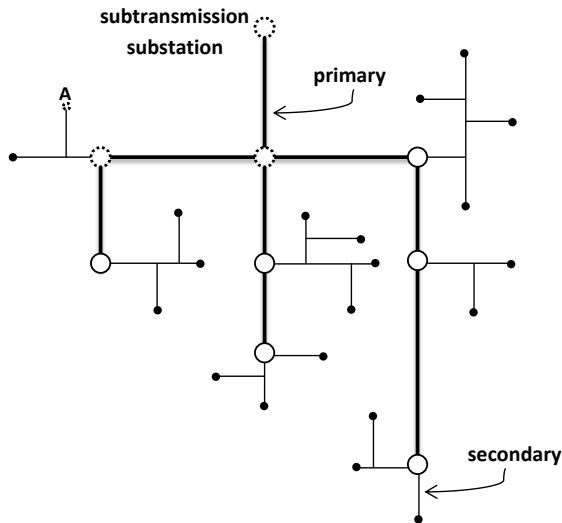


Figure 1: A one-line diagram displays the tree structure of a radial distribution network. Dotted circles represent upstream buses that supply a load connected to the pole-top transformer denoted by A.

(or buses) interconnected by lines (or feeders)<sup>1</sup>. A *one-line diagram* of a typical North American radial distribution system [16] is depicted in Figure 1. This diagram indicates each circuit irrespective of the number of phases with a single line. A load bus, represented by a circle in Figure 1, is a single point of connection on a circuit, and represents the location where power is consumed. In a distribution network a load bus is either a transformer connection to the subtransmission system (the *root* of the distribution tree), or a distribution substation from which distribution feeders originate and supply downstream distribution transformers, which step down the voltage to the secondary distribution voltage, and various loads, including homes and EV chargers. Loads are usually connected to leaf nodes of the distribution tree (*e.g.*, pole-top transformer A in Figure 1) and are not drawn in the one-line diagram.

### 3.1 Measurement and Communication

Currently, only a few measurement nodes are installed at buses. In the smart grid, we assume that many more such measurement nodes will be installed at load buses and pole-top transformers, as shown in Figure 2. A measurement node could be a current transformer which measures the current flowing through a line, a voltage transformer which measures the line voltage, or a sensor which measures the winding temperature of a transformer. These nodes are capable of continuously measuring the parameter of interest and computing an average every few milliseconds. This permits us to compute the *congestion state* of a line or a transformer. We define the congestion state as the difference between the *nominal setpoint* of the line or the transformer and its loading level. The nominal setpoint should be chosen such that continuous loading of a line or a transformer at this level

<sup>1</sup>Most distribution systems are radial; in cases where the network topology is a mesh, normally-open switches ensure that power flows only on a radial sub-graph. Thus, in this paper we assume that the distribution network forms a tree.

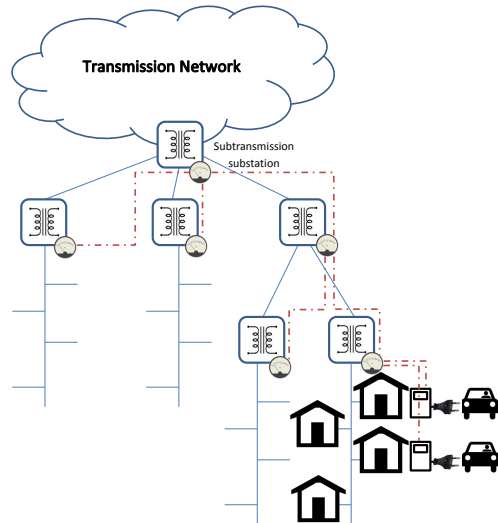


Figure 2: An illustration of a smart distribution network consisting of MCC nodes, and communication links (dashed lines).

does not cause any damage; thus, the nominal setpoint is always less than or equal to the nameplate rating. The congestion state is used to compute the *congestion price* of a line or a transformer (see Section 5). Congestion prices are used by EV chargers to set their rates as described in Section 6. They are communicated to EV chargers in a signalling cycle as explained in Section 3.3.

We distinguish between two types of congestion: *line* congestion and *transformer* congestion. A line is said to be congested when the current passing through it exceeds its rating. Likewise, a transformer is said to be congested when it is loaded higher than its nameplate rating.

In a distribution network, the protection system consisting of fuses, relays, and circuit breakers, disconnects the load shortly after a persistent line or transformer congestion is detected. We define *persistent* congestion as a congestion event which lasts for a certain number of cycles of AC power; this is a characteristic of the distribution network’s protection system. When persistent congestion is detected, a protective relay initiates the tripping action after a certain delay, which is an operating characteristic of the relay, to avoid further damage to equipment. This results in a power outage in the area supplied by the congested line/transformer. Therefore, we need to allocate charging rates to EV chargers without triggering the protection system (*i.e.*, without causing persistent congestion).

We assume that each measurement device is supplemented with a communication and control module. We refer to the entire device as a ‘Measurement, Communication, and Control (MCC)’ node. These MCC nodes form a logical tree (Figure 2). In a smart grid, a communication network will connect these MCC nodes to EV chargers (or other controllable loads) to enable the transmission of control signals.

### 3.2 Assumptions

We now lay out the assumptions that we make in this paper.

- Line or transformer overloading can be attributed en-




Figure 3: The proposed control algorithm adapts charging rates of EV chargers to the available capacity of the network.

tirely to its downstream loads. Therefore, if a measurement at a node indicates congestion, it can be ameliorated by sending a congestion signal to the node’s descendants.

- An overload can be detected within a few cycles of its occurrence, before the invocation of protection mechanisms.
- It is not possible to infer congestion implicitly at the end nodes<sup>2</sup>. Therefore, congestion must be explicitly signalled.
- The distribution system can be subject to a transient congestion for a short durations of up to  $\tau^c$  seconds before protection mechanisms, such as protective relays are brought to bear<sup>3</sup>. This gives a congestion signalling mechanism breathing room to reduce transient EV overloads. Moreover, we can choose system parameters conservatively such that an overshoot above a nominal setpoint, which is always lower than the nameplate rating, triggers congestion signalling, but does not trigger the protection system.
- The communication network is broadband, reliable, and has a low latency. We also assume that all EV chargers experience roughly the same delay when they receive signalling packets from MCC nodes (this is never exactly true but simplifies the analysis). We denote the feedback delay by  $d$ .
- The speed-of-light propagation delay between any charger and its connected substation is small, and typically less than 1ms. Thus, it is reasonable to ignore it in our model.

<sup>2</sup>In recent work, it has been shown that local sensing of the line voltage or frequency at end nodes can be used to implicitly infer the aggregate demand or the power imbalance at higher levels in the distribution network [9]. Developments arising from this pioneering work may allow implicit congestion sensing even in the electrical grid.

<sup>3</sup>We note that  $\tau^c$  is typically inversely proportional to the magnitude of the overload current.

- EV chargers are owned by electric utilities. Furthermore, they are tamper-resistant and always set their rate based on signals they receive from MCC nodes.
- An EV battery can be charged at any rate less than the maximum Amperage rating of its charger, independent of its state of charge<sup>4</sup>.
- The load from an EV charger can be increased or decreased at a fast time scale (on the order of milliseconds) with negligible effect on the EV battery lifetime.
- Rate updates at EV chargers are synchronized. Measurements of the congestion state at MCC nodes are also synchronized. This can be achieved by a broadcast time signal.
- The timescale of changes of uncontrolled loads and arrival and departure of EVs is slower than the timescale of rate updates in our distributed algorithm (Figure 3), and hence we can decouple them. This permits us to study our control problem using a model that describes a snapshot of the system in which uncontrolled loads are constant and the number of plugged-in EVs is fixed.

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. With our proposed approach, if an EV is charging at a rate that overloads the distribution system, its rate can be decreased before  $\tau^c$ , averting damage and invocation of grid self-protection mechanisms.

### 3.3 System Operation

We now describe the operation of our system. Every  $T_c$  milliseconds, the root MCC node initiates a signalling cycle by sending its congestion price to its direct children. Upon receiving the congestion price(s), an intermediate MCC node sends its own congestion price, computed using its latest recorded congestion state (as discussed in Section 5.3.1), along with the received price(s) to its children. A signalling cycle ends when EV chargers receive the congestion prices from all their parents. In Section 5.3.2, we explain how EV chargers use congestion prices of MCC nodes located on their path to set their charging rate.

## 4. OPTIMIZATION PROBLEM

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

### 4.1 Utility Function

We attribute a utility, *i.e.*, a measure of satisfaction, to an EV owner whose EV is connected to a charger denoted by  $s$ . This charger is capable of charging EVs at rates that are in  $[0, m_s]$ , where  $m_s$  is the peak charging rate that it supports. Since the departure of EVs from homes and charging stations is non-deterministic, it is reasonable to assert that EV owners are greedy and prefer to finish charging their EVs as

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

soon as possible to avoid *range anxiety*. Hence, we define the utility of a greedy EV owner by its charging rate  $x_s$ , *i.e.*, the EV owner's satisfaction is proportional to the rate at which their EV is charged.

## 4.2 Primal Problem

Our objective is to allocate the available capacity of the network fairly among EV chargers without straining the distribution network; this is generally done by solving a constrained optimization problem. There exist different fairness criteria as discussed by [11], and they are differentiated by the global objective function of the optimization problem. We adopt the notion of proportional fairness which satisfies axioms of fairness specified in game theory<sup>5</sup>. A proportionally fair allocation is in fact a Nash bargaining (arbitrated) solution [23]. It can be shown that proportional fairness is achieved if we maximize the value of a global objective function which is the sum of the logarithm of the utility functions. For notational simplicity, we denote  $\log(x_s)$  by  $U_s$ . Observe that  $U_s$  is infinitely differentiable, increasing, and strictly concave on its interior domain.

Our optimization problem would therefore become a maximization of the sum of the  $U_s$  of those chargers that are charging an EV, subject to physical constraints imposed by chargers, lines, and transformers. Since lines and transformers supply the aggregate load imposed by both homes and EV chargers, and the home loads are supposedly uncontrolled, it is necessary to subtract the load of every home from the rated capacity of its feeding transformers and lines to obtain the available capacity of every line and transformer. In effect, based on the last assumption of Section 3.2, we consider a snapshot of the system in which home loads are constant and a fixed number of EVs are plugged in to chargers.

The control problem can then be formulated as follows:

$$\begin{aligned} \max_x \quad & \sum_{s \in \mathcal{S}} U_s(x_s) & (1) \\ \text{subject to} \quad & 0 \leq x_s \leq m_s \quad \forall s \in \mathcal{S} \\ & \sum_{s: R_{sl}=1} x_s \leq c_l \quad \forall l \in \mathcal{L}, \end{aligned}$$

where  $\mathcal{S}$  and  $\mathcal{L}$  are the sets of active EV chargers, and distribution lines and transformer respectively,  $c_l$  is the available capacity of line or transformer  $l$ , and  $R$  is a  $|\mathcal{S}|$  by  $|\mathcal{L}|$  matrix encoding the topology of the network

$$R_{sl} = \begin{cases} 1 & \text{if line/transformer } l \text{ supplies charger } s \\ 0 & \text{otherwise} \end{cases}$$

We say that a line or a transformer supplies an EV charger when it is located on the path from the subtransmission substation to that charger.

This problem is a convex optimization problem as it maximizes an objective function which is the sum of concave functions (and is therefore concave), and each constraint defines a convex set. We denote a rate allocation, *i.e.*, a vector of charging rates of all chargers, by  $x = \langle x_1, \dots, x_{|\mathcal{S}|} \rangle$ . A rate allocation is *feasible* if it satisfies all constraints of the optimization problem.

<sup>5</sup>We have chosen proportionally fair resource allocation since it is the only one that provides a scale invariant Pareto optimal solution.

The two constraints of the above optimization problem can be written in matrix form:

$$\begin{aligned} 0 \preceq x \preceq m \\ xR \preceq c, \end{aligned}$$

where  $x, c$ , and  $m$  are vectors with  $|\mathcal{S}|$ ,  $|\mathcal{L}|$ , and  $|\mathcal{S}|$  components respectively, and  $\preceq$  is the generalized inequality for vectors. We refer to the second constraint as the *coupling constraint*; it couples charging rates of different EV chargers supplied by the same line or transformer.

In the next section we write the dual problem and apply the dual decomposition method to obtain a set of distributively solvable subproblems. We then design a distributed algorithm which solves the dual problem by solving these subproblems locally and independently.

## 5. CONTROLLER DESIGN

The centralized optimization problem formulated in the previous section can be solved to find a rate allocation that is proportionally fair. This requires full knowledge of the topology of the distribution network, the available capacity of lines and transformers, and the number and the location of plugged-in EVs. The distributed approach to solve the optimal control problem has three key advantages over the centralized approach. First, it gives autonomy to local controllers thereby increasing robustness of the control system. Second, it significantly reduces the communication overhead and is therefore more scalable. Third, it decreases the overall latency of control as 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 distributively solvable 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 feedback.

### 5.1 Dual Problem

The Lagrangian function of our optimization problem is

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

where  $\lambda = (\lambda_1, \dots, \lambda_{|\mathcal{L}|})$  is a vector of Lagrangian multipliers associated with the coupling constraints, and

$$y_l = \sum_{s: R_{sl}=1} x_s \quad \forall l \in \mathcal{L}$$

Thus, the Lagrangian dual problem would be

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

which is equivalent to

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

where

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

In the above equation,  $f(x; \lambda)$  represents  $f$  as a function of  $x$  parameterized by  $\lambda$ . 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.

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

An important remark is that strong duality holds as all inequality constraints are affine. Therefore, we can write the following KKT optimality conditions

$$\frac{1}{\hat{x}_s} = \sum_{l:R_{sl}=1} \hat{\lambda}_l \quad \forall s \in \mathcal{S} \quad (6)$$

$$0 \leq \hat{x}_s \leq m_s \quad \forall s \in \mathcal{S} \quad (7)$$

$$\hat{\lambda}_l \geq 0 \quad \forall l \in \mathcal{L} \quad (8)$$

$$\hat{\lambda}_l (\hat{y}_l - c_l) = 0 \quad \forall l \in \mathcal{L} \quad (9)$$

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 last 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  $\hat{\lambda}$ .

$$\hat{x}_s = \min \left\{ \frac{1}{\sum_{l:R_{sl}=1} \hat{\lambda}_l}, m_s \right\} \quad (10)$$

## 5.2 Dual Decomposition

Writing the Lagrangian dual problem in the form of (4) reveals its hidden decomposition structure [18]. Particularly, each EV charger can locally solve a subproblem given by

$$\max_{0 \leq x_s \leq m_s} f_s(x_s; \lambda), \quad (11)$$

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* [12]) 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

$$\min_{\lambda \geq 0} \left\{ \sum_{l \in \mathcal{L}} \lambda_l c_l + \sum_{s \in \mathcal{S}} f_s(\hat{x}_s; \lambda) \right\}. \quad (12)$$

where  $f_s(\hat{x}_s; \lambda)$  is the optimal value of (11). Observe that the objective function of the master problem is linear in  $\lambda$  and its derivative with respect to a Lagrangian multiplier is given by

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

## 5.3 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 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 6. In Section 7, we specify sufficient conditions for convergence of this algorithm to primal and dual optimal values.

### 5.3.1 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 (12). The following algorithm updates congestion prices in each iteration in opposite direction to the gradient of the dual function.

$$\lambda_l(t+1) = \max\{\lambda_l(t) - \kappa(c_l - y_l(t)), 0\} \quad \forall l \in \mathcal{L} \quad (13)$$

Here  $\kappa$  is a sufficiently small positive constant which determines the responsiveness and stability of control, and has to be selected carefully. Note that it is not necessary to estimate  $c_l$  and  $y_l$  at an MCC node to compute  $c_l - y_l$ . This is because  $c_l - y_l$  is equal to the line or transformer loading subtracted from its nominal setpoint denoted  $\xi_l$ . The nominal setpoint of a line or transformer is known a priori and its loading can be measured by the corresponding MCC node.

### 5.3.2 A Control Law for Adjusting the Charging Rate

We denote the latest congestion price vector that an EV charger received by  $\lambda(\bar{t})$ . The subproblem (11) can be easily solved by finding the stationary point of  $f_s(x_s; \lambda)$ .

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

$$x_s(t) = \min \left\{ \frac{1}{\sum_{l:R_{sl}=1} \lambda_l(\bar{t})}, m_s \right\} \quad (14)$$

Note that  $x_s(t)$  would be 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<sup>6</sup>. More specifically,  $y_l(t)$  is given by

$$y_l(t) = \sum_{s:R_{sl}=1} x_s(t) \quad \forall l \in \mathcal{L} \quad (15)$$

We end this section by a remark that the unit of time in (13) and (14) is  $T_c$  milliseconds, and therefore  $\bar{t}$  equals  $t - \frac{d}{T_c}$  as congestion prices are received by EV chargers after  $d$  milliseconds.

<sup>6</sup>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.

## 6. 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 5.

### 6.1 Normal Operation

During normal operation, our distributed charging control algorithm measures the congestion state of a line or a transformer and computes the corresponding congestion price based on (13). This price is sent to descendant EV chargers in a signalling cycle initiated by the root MCC node every  $T_c$  milliseconds. Recall we assume that MCC nodes are synchronized and therefore update their congestion prices at the same time (see Algorithm 1).

---

**Algorithm 1:** Congestion price update at MCC node  $l$  with nominal setpoint  $\xi_l$

---

**input:**  $\xi_l, \kappa (> 0)$

**while true do**

- Measure load
- congestion state  $\leftarrow \xi_l - \text{load}$
- price  $\leftarrow \max\{\text{price} - \kappa \times \text{congestion state}, 0\}$
- Send price along with all received prices to children
- Wait until the next **clock tick**

**end**

---

After receiving congestion prices from upstream MCC nodes, every charger computes its charging rate using (14) and starts charging at this rate. We assume that EV chargers are also synchronized, i.e., they adjust their rates at the same time.

---

**Algorithm 2:** Rate adjustment at EV charger  $s$

---

**input:**  $m_s$ , new congestion prices

**while true do**

- $\lambda \leftarrow \text{vector of new congestion prices}$
- aggregate price  $\leftarrow \sum_{l \in \text{ascendants}} \lambda_l$
- rate  $\leftarrow \min\{\frac{1}{\text{aggregate price}}, m_s\}$
- Start charging the battery at **rate**
- Wait until the next **clock tick**

**end**

---

Note that the clock ticks every  $T_c$  in both algorithms.

### 6.2 Emergency Response Mode

To assure grid reliability at all times, we have designed a ‘fail-safe’ load reduction mechanism to respond to sudden load spikes. This mechanism operates both at the MCC nodes and at EV chargers. If any MCC node detects that its overall load has exceeded the nominal setpoint of  $\xi_l$  by a factor  $\eta$ ,  $\eta > 1$  for a duration that exceeds  $\tau^c - d$ , it sends an emergency shutdown signal to all descendant EV chargers. In response, all EV chargers reduce their charging rate to 0. Since this response is guaranteed to move the system from an overloaded state to an underloaded state within time  $\tau^c$ , it avoids triggering the protection system, and at the same time does not affect our proofs of system stability and convergence, discussed next.

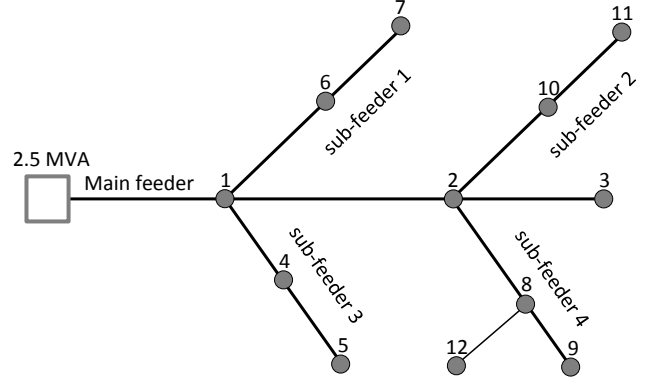


Figure 4: The one-line diagram of our test distribution network.

## 7. CONVERGENCE ANALYSIS

This section studies the conditions under which the proposed distributed control algorithm converges to the solution of the centralized optimization problem (1) in a static context, i.e., no EVs arrive or depart and the change in the magnitude of uncontrollable loads is negligible. Recall that the primal optimum is equal to the dual optimum as strong duality holds. Therefore, we instead show that the distributed control algorithm converges to the solution of (3).

Our proof technique is to extend Theorem 1 in [14] to derive the following result.

*Theorem 1.* Starting from any initial rates  $0 \preceq x \preceq m_s$  and congestion prices  $\lambda \succeq 0$ , the distributed algorithm introduced in Section 6 converges to the primal-dual optimal values if

- (1)  $T_c \geq d$
- (2)  $0 < \kappa < \frac{2}{\bar{m}^2 \bar{L} \bar{S}}$

where  $\bar{m} := \max_s m_s$ ,  $\bar{L} := \max_s \sum_l R_{sl}$ , and  $\bar{S} := \max_l \sum_s R_{sl}$ .

Recall that  $\bar{L}$  is the maximum number of lines and transformers which are instrumented with MCC nodes and are supplying the load of an EV charger,  $\bar{S}$  is the total number of EVs being charged in the distribution network, and  $\bar{m}$  is the maximum charging rate supported by an EV charger.

To prove this theorem, we extend the proof in [14] by showing that the first condition is satisfied when the communication delay is bounded by  $d$ . Recall that our algorithm requires MCC nodes to update congestion prices periodically; thus, congestion prices are unchanged between two consecutive updates which are  $T_c$  milliseconds apart. Assuming that  $t - \frac{d}{T_c}$  is the last time that congestion prices are updated, we have  $\lambda(t - \frac{d}{T_c}) = \lambda(t)$  if  $T_c \geq d$ . Substituting  $\lambda(t - \frac{d}{T_c})$  with  $\lambda(t)$  in (14), our algorithm becomes identical to the one presented in [14].

## 8. RESULTS

The behavior of our control algorithm is determined by the value of two critical parameters  $\kappa$  and  $T_c$ . The first parameter determines the step size in changing congestion prices: the larger this value, the more responsive the system. The




Figure 5: The value of the step size  $\kappa$  determines how the loading of the substation transformer changes over time (Scenario A).

second parameter determines the frequency of control: the larger this value, the less the control overhead, but the longer the time that the system may be experiencing transient congestion (to prevent triggering the protection system, we require  $T_c < \tau^c - d$ ). Theorem 1 states sufficient conditions under which the proposed distributed algorithm converges. However, the space of possible parameters defined by Theorem 1 is large and depending on their values the control system exhibits different characteristics. In this section, we simulate our algorithm in a test distribution network to investigate the sensitivity of the control to the choice of these two parameters, and decide on their optimal values considering overall system constraints.

## 8.1 Test Distribution System

Our analysis is based on a 4.16kV radial distribution system comprised of 13 buses. This system is obtained by simplifying the IEEE 13-bus test feeder [1]. In particular, we do not model in-line transformers, voltage regulators, shunt capacitor banks, and protection switches. Moreover, we assume that the three phase system is balanced, *i.e.*, loads are distributed as evenly as possible between the phases. Thus, we perform a *per-phase analysis* of the network, assuming that all lines and transformers are single-phase.

Figure 4 depicts our test distribution network consisting of a main feeder connecting buses 1, 2, and 3 (the circles labelled with these numbers in the figure) to the substation and four sub-feeders branching out from the main feeder. The main feeder is rated at 730A and sub-feeders 1 to 4 are rated at 340A, 340A, 140A, and 230A respectively. The network is supplied by a single phase 2.5MVA transformer which is installed at the substation and steps down the voltage from 115kV to 4.16kV. The voltage is further reduced to the secondary distribution voltage by pole-top transformers that are connected to buses. For sake of simplicity, in this paper we do not model these transformers and feeders radiating from them, although the analysis would be exactly similar to what we do in this section.




Figure 6: The number of iterations it takes to achieve 95% convergence (Scenario A).

We assume that the nominal setpoint of each line or transformer is equal to its nameplate rating, and every bus in the test distribution network is instrumented with an MCC node. Spot loads are connected to selected load buses. We assume that they draw constant current from the network and the power factor at the substation bus is one. This permits us to ignore the reactive power flow in our simulations. Furthermore, it is assumed that EV chargers that support charging at up to 16A are also connected to selected load buses.

Table 1 summarizes our simulation scenarios. A simulation scenario describes a snapshot of the system in which constant current spot loads are connected to buses, and a number of EVs are charged by chargers which are downstream of these buses. We choose Scenario A so that the substation transformer is the only congested resource when allocating proportionally fair rates to EV chargers. In contrast, in Scenario B, both the substation transformer and the line connecting buses 2 and 8 become congested simultaneously when allocating proportionally fair rates to EV chargers.

## 8.2 Sensitivity Analysis

In this section, we compare the number of iterations that it takes to achieve 95% convergence and the maximum deviation from the nominal setpoint value of a line or a transformer for different values of  $\kappa$ .

Consider the first simulation scenario. In this case, the maximum step size for which the convergence of the algorithm is guaranteed is  $\kappa^* = \frac{2}{16^2 \times 18 \times 5} = 8.68 \times 10^{-5}$ . As we increase the value of  $\kappa$ , the control system transitions from an over-damped system to an under-damped system and eventually to an unstable system for  $\kappa > 8.05 \times 10^{-4}$  which is much larger than  $\kappa^*$ . Figure 5 shows how the loading of the substation transformer varies over different iterations for three different values of  $\kappa$ .

Changing the value of  $\kappa$  changes the system responsiveness, specifically, the number of iterations it takes to achieve the 95% convergence, where the 95% convergence is said to



The bus number	3		4		5		6		7		8		9		10		11		12	
Scenario A	30	2	40	2	40	2	10	2	20	2	70	0	20	2	100	2	20	2	40	2
Scenario B	30	2	40	2	40	2	10	2	20	2	160	0	20	2	10	2	20	2	40	2

Table 1: For each bus the first column shows the spot load connected to it (in Ampere) and the second column shows the number of EV chargers supplied by it.




Figure 7: The maximum deviation from the rating of the substation transformer (Scenario A).

be achieved when the loading of a line or transformer remains in the boundary given by  $\pm 5\%$  of its nominal set-point. Figure 6 shows that the number of iterations required for 95% convergence decays exponentially as we increase the value of  $\kappa$ . When the step size is equal to  $\kappa^*$ , it takes 242 iterations to achieve 95% convergence; assuming that  $T_c$  is 20 milliseconds the algorithm achieves 95% convergence in 4.8 seconds. Should we set the step size larger, for example  $\kappa = 4 \times 10^{-4}$ , it takes only 54 iterations to achieve 95% convergence which is equivalent to approximately one second (though, at this value of  $\kappa$ , convergence is not guaranteed!).

We find that the maximum deviation from the set point value is almost 4% for a wide range of step size values including  $\kappa^*$  (see Figure 7). However, once  $\kappa$  is large, the system starts oscillating around the set point value and the maximum deviation could be relatively high.

Now consider the second simulation scenario. We find that similar to Scenario A, the number of iterations it takes to achieve 95% convergence decays exponentially as we increase the value of  $\kappa$ . If we set the step size to  $\kappa^*$ , the algorithm converges after 354 iterations and the maximum deviation is about 1%.

Lastly, we validate the assumption that if the timescale of changes of uncontrolled loads and arrival and departure of EVs is slower than the timescale of rate updates in our algorithm, the stability result obtained for a static network scenario can be extended to the dynamic case. To see this, suppose that uncontrolled loads vary every 5 seconds in our test distribution network, and the system can be described in an interval of 15 seconds by the following sequence of snapshots: Scenario A, Scenario B, and again Scenario A. Figure 8 and




Figure 8: Loading of the substation transformer as the system transitions from Scenario A to Scenario B after 250 iterations and back to Scenario A again after another 250 iterations.

Figure 9 show the substation transformer loading and the loading of the line connecting buses 2 and 8 respectively, assuming that the step size is  $\kappa^*$ , and  $T_c$  is 20 milliseconds. It can be readily seen that since  $\kappa$  and  $T_c$  are chosen according to Theorem 1, transitions from one scenario to another one does not destabilize the system.

As an extreme case, we consider the situation where the home load goes from zero (so that all EVs are charged at their maximum rate) to Scenario A, where the transformer is congested, and back to zero in the three time periods (Figure 10). Here, at time 5s, without the emergency response mechanism, the transient congestion at the transformer lasts for 30 iterations or 0.6s. However, by adding the emergency response mode which is invoked after  $5T_c = 100\text{ms}$ , the transient load does not exceed the design specification of  $\tau^c = 200\text{ms}$ , avoiding invocation of the grid protection system.

### 8.3 Proportional Fairness

The solution to the centralized optimization problem (3) indicates that the fair share of every EV charger in Scenario A must be 11.72A. Our results also show that when the algorithm converges, 11.72A is allocated to each charger. Figure 11 shows the charging rates of a few EV chargers over time.

In Scenario B, the proportionally fair rate of chargers connected to buses 9 and 12 is 2.5A, whereas the proportionally fair rate of all other chargers is 14.35A (see Figure 12). This is because chargers connected to buses 9 and 12 have both bottleneck resources on their path to the substation, and therefore their fair rate would be smaller than other charg-




Figure 9: Loading of the line connecting buses 2 and 8 as the system transitions from Scenario A to Scenario B after 250 iterations and back to Scenario A again after another 250 iterations.

ers. Note that the optimal congestion price of non bottleneck lines and transformers is zero (see Equation 9).

#### 8.4 Design Space

As discussed earlier, our design space is defined by the conditions of Theorem 1. Clearly, the step size should be equal to  $\kappa^*$  to increase responsiveness of the control system, thereby reducing the number of iterations it takes to achieve 95% convergence. Once we fix the value of  $\kappa$ , we decide on the value of  $T_c$ . Since the 95% convergence time is proportional to  $T_c$ , we prefer to set it as small as possible, *i.e.* equal to  $d$ .

However, in practice  $T_c$  is also bounded from below by another parameter. It turns out that sending congestion prices periodically from MCC nodes to EV chargers causes a significant communication overhead. In effect, the communication medium and the protocol chosen for transmission of these signals determine the fastest rate at which we can propagate updates in our algorithms. Hence, although the communication delay is of the order of a few milliseconds, it is not feasible to update congestion prices faster than every 20 milliseconds.

### 9. DISCUSSION AND CONCLUSION

Our work represents a novel approach to controlling the charging of electric vehicles. Instead of forecasting the number of EVs and the non-EV load several hours ahead and solving an optimization problem, we rely on fast measurements and communication to avert persistent congestion and the invocation of protection mechanisms. Moreover, by using a mathematical framework originally developed for rate control in the Internet, each EV charger can independently update its charging rate, yet the global allocation converges to the nominal operating setpoint, the allocated rates are proportionally fair, and the allocation is optimal. Additionally, the use of an emergency response mode averts protection events without compromising these benefits.




Figure 10: Loading of the substation transformer as the system transitions from zero home load to Scenario A after 250 iterations and back to zero home load again after another 250 iterations.

Our control algorithm can be used to implement demand response (via EV charging) in a distribution network. This requires defining a new type of congestion which corresponds to generation shortfall. Concretely, when supply falls short of demand, the system operator sends a signal to root MCC nodes. Consequently, each root MCC node indicates congestion by increasing its congestion price or sending an emergency shutdown signal to controlled loads in its distribution tree.

A possible extension to our work would be to investigate the voltage drop problem at distant buses. This requires adding other constraints to the centralized optimization problem. A voltage drop can be modeled as line congestion. Corresponding MCC nodes should update their congestion state based on the difference between the measured line voltage and their rated voltage.

Another direction for future work is to design an asynchronous distributed algorithm to control EV charging, similar to the algorithm proposed in [14]. This eliminates the need for time synchronization between EV chargers and MCC nodes and would give us a new bound on step size values for which convergence of the algorithm is guaranteed.

It is interesting to compare our distributed approach with a centralized one. They can be compared in terms of robustness, efficiency, communication overhead, and control and communication delay. The distributed approach is more robust as there is no single point of failure. The centralized approach is potentially more efficient if the primal optimization problem can be solved quickly. Both approaches are expected to have roughly the same overhead and delay. A further comparison of these two approaches is left for future work.

The main limitation of our approach is that it requires a heavy communication overhead; our design point is one message to each EV every 20ms. Reducing this overhead by decreasing the communication frequency, or by using an un-




Figure 12: Changes to charging rates of three chargers over time obtained for Scenario B.

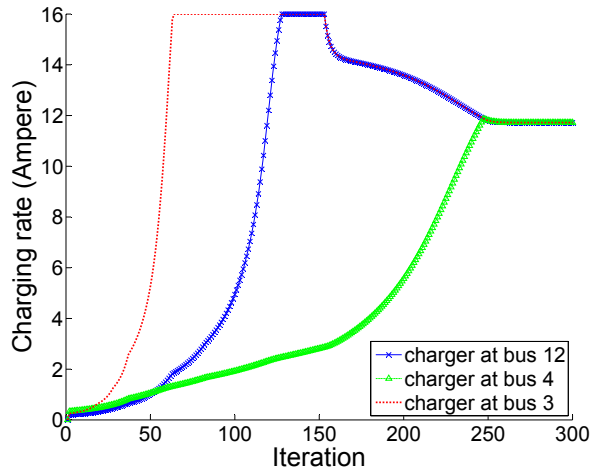


Figure 11: Changes to charging rates of three chargers over time for Scenario A.

derlying broadcast medium, such as 700 MHz cellular radio, is a clear direction for future work.

A second problem with our approach is that it does not scale well: the choice of  $\kappa$ , which controls system responsiveness, is constrained by the number of EVs and the maximum EV charging rate. As these increase,  $\kappa$  and system responsiveness decrease. What is needed is a less conservative bound for  $\kappa$  that has better scaling properties. Alternatively, we could use a larger value of  $\kappa$  than  $\kappa^*$  and then rely on the emergency response mode to address any resultant overloads. We plan to investigate these alternatives in future work.

Finally, our operating parameters, such as the protection threshold  $\tau^c$ , safety margin  $\eta$ , and communication interval  $T_c$  would be better chosen from a realistic EV charging testbed. We hope to validate our choice of these parameters through a deeper engagement with electric utilities.

## 10. REFERENCES

- [1] IEEE Distribution Test Feeders. <http://ewh.ieee.org/soc/pes/dsacom/testfeeders>.
- [2] SAE J1772 Standard. [http://standards.sae.org/j1772\\_201210](http://standards.sae.org/j1772_201210).
- [3] C. Ahn, C.-T. Li, and H. Peng. Optimal decentralized charging control algorithm for electrified vehicles connected to smart grid. *J. Power Sources*, 196(2):10369 – 10379, 2011.
- [4] O. Ardakanian, C. Rosenberg, and S. Keshav. Real-time distributed congestion control for electrical vehicle charging. *Performance Evaluation Review*, 2012.
- [5] A. G. Boulanger, A. C. Chu, S. Maxx, and D. L. Waltz. Vehicle Electrification: Status and Issues. *Proceedings of the IEEE*, 99(6):1–23, 2011.
- [6] R. Brown. Impact of smart grid on distribution system design. In *IEEE PES General Meeting*, pages 1–4. IEEE, 2008.
- [7] EPRI. Impact of Plug-in Hybrid Electric Vehicles on Utility Distribution . Technical Report, 2009.
- [8] L. Gan, U. Topcu, and S. Low. Stochastic distributed protocol for electric vehicle charging with discrete charging rate. Technical report, California Institute of Technology, 2011.
- [9] T. Ganu, J. Hazra, D. P. Seetharam, S. A. Husain, V. Arya, L. C. De Silva, R. Kunnath, and S. Kalyanaraman. nPlug: a smart plug for alleviating peak loads. In *ACM e-Energy*, 2012.
- [10] X. Gong, T. Lin, and B. Su. Survey on the impact of electric vehicles on power distribution grid. In *Power Engineering and Automation Conference*, volume 2, pages 553–557. IEEE, 2011.
- [11] F. Kelly. Charging and rate control for elastic traffic. *European transactions on Telecommunications*, 8(1):33–37, 1997.
- [12] 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*, 49(3):237–252, 1998.
- [13] J. Lopes, F. Soares, and P. Almeida. Integration of electric vehicles in the electric power system. *Proceedings of the IEEE*, 99(1):168–183, 2011.
- [14] S. H. Low and D. E. Lapsley. Optimization flow control. I. Basic algorithm and convergence. *Networking, IEEE/ACM Transactions on*, 7(6):861–874, 1999.
- [15] Z. Ma, D. S. Callaway, and I. A. Hiskens. Decentralized charging control of large populations of plug-in electric vehicles. *IEEE Trans. Control Systems Technology*, (99):1–12, 2011.
- [16] A. Meier. *Electric Power Systems: A Conceptual Introduction*. Wiley-IEEE Press, 2006.
- [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*, 13(1):43–56, 2005.
- [18] D. Palomar and M. Chiang. A tutorial on decomposition methods for network utility

- maximization. *Selected Areas in Communications, IEEE Journal on*, 24(8):1439–1451, 2006.
- [19] S. Shenker. Fundamental design issues for the future internet. *IEEE Journal on Selected Areas in Communications*, 13:1176–1188, 1995.
- [20] E. Sortomme, M. Hindi, S. MacPherson, and S. Venkata. Coordinated charging of plug-in hybrid electric vehicles to minimize distribution system losses. *IEEE Trans. Smart Grid*, 2(1):198–205, 2011.
- [21] R. Srikant. *The Mathematics of Internet Congestion Control (Systems and Control: Foundations and Applications)*. Birkhauser, 2004.
- [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](http://www.cisco.com/web/strategy/docs/energy/control_architecture.pdf).
- [23] H. Yaïche, R. R. Mazumdar, and C. Rosenberg. A game theoretic framework for bandwidth allocation and pricing in broadband networks. *IEEE/ACM Trans. Networking*, 8(5):667–678, 2000.
- [24] L. Ying, G. Dullerud, and R. Srikant. Global stability of internet congestion controllers with heterogeneous delays. *Networking, IEEE/ACM Transactions on*, 14(3):579–590, 2006.