Real-Time Distributed Congestion Control for Electrical Vehicle Charging

(Invited Paper)

O. Ardakanian	C. Rosenberg	S. Keshav
David R. Cheriton School of Computer Science	Electrical and Computer Engineering Dept.	David R. Cheriton School of Computer Science
University of Waterloo	University of Waterloo	University of Waterloo
oardakan@uwaterloo.ca	cath@uwaterloo.ca	keshav@uwaterloo.ca

ABSTRACT

The significant load and unpredictable mobility of electric vehicles (EVs) makes them a challenge for grid distribution systems. Unlike most current approaches to control EV charging, which construct optimal charging schedules by predicting EV state of charge and future behaviour, we leverage the anticipated widespread deployment of measurement and control points to propose an alternative vision. In our approach, drawing from a comparative analysis of Internet and distribution grid congestion, control actions taken by a charger every few milliseconds in response to congestion signals allow it to rapidly reduce its charging rate to avoid grid congestion. We sketch three control schemes that embody this vision and compare their relative merits and demerits.

1. INTRODUCTION

The large-scale introduction of electric vehicles (EVs) is likely to greatly affect the electrical grid's distribution system [3, 6]. This is primarily because each EV can impose a significant load on the distribution network: with Level 2 charging, EVs can be charged at up to 40A at 240V, a load of 9.6kW, 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 ten average homes! Using lower-level (i.e., Level 1) charging does reduce the impact on the grid but only at the expense of greatly increasing the duration of the charging process. Therefore, there is an inherent trade-off between charging level, charging duration, and impact on the grid.

A second challenge posed by EVs is that their load is time-variable: an EV's load *unpredictably* disappears when it is driven. EV mobility has the additional impact that EV load may appear at different parts of the distribution network at different times.

To deal with these challenges, electric utilities, who own and operate the distribution system, must either upgrade their distribution assets or intelligently control the EV charging process to manage their load. The former is unaffordably expensive; hence the great interest in the latter, which is also the focus of our work.

Existing approaches to control EV charging either use a central controller to coordinate charging [8] or cast the control algorithm in the form of a distributed optimization [1]. 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. 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, and their initial state of charge. The safety margin built in to hedge against prediction errors makes both approaches overly conservative.

The future smart grid is likely to have a considerable number of measurement and control devices that are interconnected by a ubiquitous communication network [4]. Therefore, we propose to use fast-timescale measurements and communication to control EV charging, motivated by techniques for congestion control in the Internet.

Specifically, we note that a typical EV charger is located within 300km of the nearest substation. Therefore, the speed-of-light propagation delay between any charger and its connected substation is less than 1ms. This implies that it is feasible to design and implement a control algorithm that changes the EV charging rate in response to the congestion state of the distribution system (which depends in part on the uncontrollable loads) every few milliseconds, which is on the same order of magnitude as one cycle of AC power (16.6ms). With our proposed approach, if an EV is charging at a rate that affects the reliability of the grid, that is, overheating a transformer or overloading a feeder, its rate can be decreased in a few cycles, averting damage and the invocation of grid self-protection. This fundamental insight changes the approach to EV charging from a slow central or decentralized optimization approach to a *fast dynamic* approach.

We make three specific contributions:

- We show that the congestion control problem in the context of a distribution system is similar in many aspects to the congestion control problem in the Internet
- We propose a measurement and signalling architecture to provide real-time explicit feedback to EV chargers
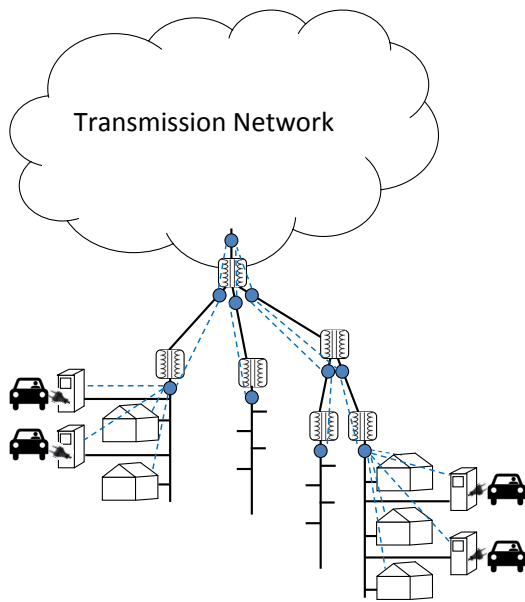
Figure 1: The physical model of a distribution system consisting of transformers, feeders (solid lines), MCC nodes (circles), and communication links (dashed lines).

- We present three real-time distributed congestion control mechanisms for charging EVs

Our focus is on establishing a vision and proposing a high-level architecture, rather than validation and the analysis, which we defer to further study.

The rest of the paper is outlined as follows. Section 2 overviews the structure and operation of a power distribution system and compares it with the Internet to show that congestion is defined similarly in both networks. Section 3 specifies our design goals and presents three congestion control mechanisms by describing the joint measurement and signalling architecture, and the decision making algorithms. We mention some discussion points and future directions and conclude the paper in Section 4.

2. BACKGROUND

2.1 Overview

The electrical grid consists of generation, transmission, and distribution systems. The distribution system, which is the focus of this paper, is responsible for delivering electricity to end customers. A radial distribution system, shown in Figure 1, can be modelled as a tree where transformers and loads (either controllable like EV chargers or uncontrollable like homes) are its vertices, and feeders are its edges[1]. The *root* of the tree is a subtransmission substation which connects the transmission network to the distribution network by stepping down the transmission voltage to primary distribution voltage. Feeders radiating from the substation bus take power to industrial consumers and downstream distribution transformers. Depending on the distance and the

---

[1]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
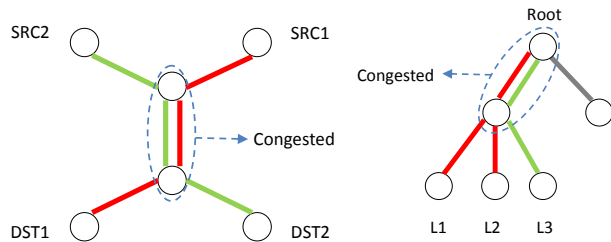

Figure 2: a) a congested link in the Internet may affect nodes sending data to different destinations. b) a congested feeder in the distribution network only affects downstream loads connected to this feeder.

load, one or more levels of distribution transformers (in addition to the subtransmission transformer) step down the voltage to the distribution voltage and supply various loads including homes and EVs chargers (*i.e.*, leaves of the tree) along the feeders. Currently, only a few measurement nodes, such as current transformers that measure electricity current, are installed at substation buses. In the future smart grid, we assume that there will many more such measurement nodes, as shown in Figure 1. We refer to a measurement node supplemented with a communication module as an MCC node (described in Section 3.4). In a smart grid, a communication network will connect these MCC nodes to enable transmission of local measurements to other parts of the grid.

We assume that EV chargers may or may not be owned by utility companies. Thus they may not be tamper-resistant and may not always conform to utility's signals. We also assume that EV chargers support variable-rate charging and this has negligible impact on the lifetime of EV batteries. Based on these assumptions, our goal is to adapt the charging rate of EVs so that the overall distribution system is both congestion-free and efficient.

2.2 Congestion Control: the Internet vs. the Grid

This section compares a packet-switched communication network such as the Internet with a distribution network. We focus on the nature of congestion control in either system, and why we believe a real-time distributed feedback control mechanism can be adopted in a distribution system. Our criteria for comparison are as follows:

- *Definition of congestion:* In communication networks, a path from a source to a destination is said to be congested when the buffer content of at least one of the routers on this path is persistently above a specific threshold. Similarly, a distribution network is said to be congested when the current passing through at least one feeder persistently exceeds its current limit or the winding hot spot temperature of at least one transformer exceeds a specific threshold (this is known as transformer thermal overloading).
- *Topology:* The Internet is a general mesh network consisting of many sources and destinations that are connected by communication links and routers. If a source congests a link by sending a burst of data, all other sources that send data through this link see its impact on their quality of service (*e.g.*, their average end-to-end delay increases) even if their packets are going to

different destinations (Figure 2(a)). A radial distribution system has a tree topology in which every node has a parent which supplies its demand and the *root* of the tree supplies the demand of all loads in this tree. If a few loads congest a feeder by consuming high power, only downstream loads that are supplied by this feeder are affected (Figure 2(b)). Other loads located in this tree will not be affected.

- *Infrastructure for sending measurement and control signals:* In the Internet, data packets carry control information and therefore the same infrastructure is used for transmitting both data and control signals. In contrast, power lines deliver electricity to customers and congestion signalling is done separately (using an auxiliary telecommunication network).

- *Congestion notification:* There are two types of congestion feedback in the Internet: explicit and implicit. Intermediate routers can *explicitly* report congestion (or the degree of congestion) to end nodes. End nodes can also infer congestion by measuring packet loss or estimating the round-trip delay; this is known as *implicit* congestion notification. However, in the grid it is almost impossible to infer congestion implicitly[2].

- *Self-protection:* Both systems protect themselves against congestion. For example, Internet routers are configured to drop packets to avoid congestion. In a distribution network, the protection system consisting of relays and circuit breakers trip and disconnect the load in case of congestion. The protection mechanisms of these two networks differ in that the packet dropping scheme implemented by routers does not interrupt service to clients (though it may impact the quality of service); however, when a protective relay trips all downstream loads are disconnected, leading to a service disruption.

- *Uncontrolled loads:* Both systems are designed to deal with uncontrollable demands. Specifically, UDP traffic is uncontrolled in the Internet and congestion control mechanisms do not deal with this type of traffic (although UDP packets can easily be filtered and discarded if necessary). Similarly, there are uncontrollable loads in the grid which do not respond to control signals. The main difference is that the current infrastructure does not permit the segregation of the uncontrolled loads from the controlled ones and *vice-versa*.

3. DESIGN OF CONGESTION CONTROL MECHANISMS

We begin this section by sketching our vision. We then specify the desired design requirements that a congestion control mechanism for charging EVs should satisfy. Finally, we describe our proposed measurement architecture, signalling infrastructure, and three decision making algorithms.

3.1 The Vision

We envision a system in which the chargers communicate with measurement nodes in the distribution system by

means of a reliable and fast communication network. When an EV needs charging, it starts charging at a low rate and then increases it slowly as long as it does not receive a signal from the grid that indicates congestion. This congestion might be due to the chargers themselves or to an increase in the uncontrollable loads. Thanks to the efficient communication and control infrastructure, the charger can react nearly immediately to congestion signals, averting the use of grid protection actions from circuit breakers. This vision is the one used in the telecommunication networks to control elastic flows (e.g. with FAST TCP) and it has proven to be both stable and efficient [11]. Importantly, with this approach, chargers can charge at high rates if the grid is underloaded. For this vision to succeed, there must be a joint measurement and signalling infrastructure in place to detect the outset of congestion very quickly and to inform the chargers that are in the congested region. Moreover, only chargers that are in the congested region should need to decrease their rates.

Note that this vision does not require us to predict EV state of charge or their mobility: charging happens for those EVs that are present in the system at any point in time, and their charging rates are controlled every few milliseconds. This allows us to move away from prediction and quickly respond to changes in available charging capacity due to fluctuations in uncontrollable loads.

3.2 Design Goals

A control mechanism for charging EVs is expected to satisfy the following properties.

- **Maintaining grid reliability**: Distribution assets, such as feeders and transformers, should never be overloaded. This averts load shedding and maintains the normal lifetime of these assets. Protective relays are designed to trip and disconnect the power supply to a neighbourhood when a feeder's current limit is exceeded or a transformer is persistently overloaded. Thus, a crucial goal in the design of a control mechanism for charging EVs is to maintain the same level of reliability and to ensure that no additional power outage is introduced due to their charging.

- **High utilization**: It is almost always favorable to increase the utilization of the distribution system, as it implies better expenditure management. The maximum utilization is achieved when the steady state (*i.e.*, long-term) loading of an asset equals its nominal rating and it takes the minimum time for the system to converge to this steady state. Because real systems rarely converge smoothly to their steady state loading, overshooting the target to some extent, we set the target steady state loading of an asset to its nominal rating times a constant factor, λ ($0 < \lambda \leq 1$). This constant factor is used to hedge against the risk of system overloading due to transient system behaviour.

- **Minimize oscillations**: Oscillations are usually inefficient. Moreover, frequent oscillations in the EV charging rate could have a negative impact on the lifetime of EV batteries; thus, they should be minimized.

- **Fairness**: Allocation of charging rates to EV chargers must be done according to a fairness criteria, for example, by offering the max-min fair share [9] to each charger.

- **Robustness**: The overall system is expected to be ro-

---

[2]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 [7]. Developments arising from this pioneering work may allow implicit congestion sensing even in the electrical grid.

bust against failure of the signalling network.

3.3 Measurement

The distribution network is assumed to be equipped with measurement devices as shown in Figure 1: A current transformer installed at the point of connection of each feeder to the substation bus continuously measures the current in the feeder and computes an average every t_M ms. A thermometer installed at each substation continuously measures winding hot spot temperature of the transformer and computes an average every t_M ms. This allows us to compute the difference between the current limit[3] of the feeder and the current passing through it as well as the difference between the maximum winding temperature and the measured winding temperature of the transformer. If these differences are positive it means that the feeder/transformer is not overloaded. Otherwise, it is nearly overloaded (depending on the value of λ) and the protective relay will, most probably, trip if this condition persists or worsens.

3.4 Signalling and Communication

Our congestion control mechanisms assume that each measurement device is supplemented with a communication and control module. We refer to the entire device as an MCC node. A communication network interconnects these MCC nodes in the distribution network as illustrated in Figure 3. Specifically, every MCC node, except the root MCC node that is installed at the primary side of the subtransmission substation, is connected to its parent, which is the MCC node installed at the upstream feeder, and its children, which are either the MCC nodes installed at downstream feeders or at controllable loads supplied by the corresponding feeder. The root MCC node does not have a parent thus, it is only connected to its children. We assume that the communication network is reliable and has a low latency, and all EV chargers experience roughly the same round-trip-time when they send or receive a packet to the root MCC node.

Figure 3 is an abstraction of the underlying joint measurement and signalling infrastructure which merely includes MCC nodes and controllable loads. The feedback on the congestion status is propagated hop-by-hop until it reaches the destination.

We distinguish between two possible types of explicit feedback: one that contains information about the degree of congestion (*e.g.*, the difference between the current limit of the feeder and the measured current passing through it) and the one that simply indicates the existence of congestion (*e.g.*, a boolean that tells whether a feeder/transformer is congested or not).

3.5 Decision Making

The decision on how to reduce a charger's load on receiving a grid congestion signal is called the *distributed rate allocation* algorithm. The space of possible rate allocation algorithms is large, just as the space of Internet congestion control schemes is large (see [5] for a survey of many Internet

---

[3]We define the current limit of the feeder as the rated maximum current that can pass through it multiplied by a constant factor, λ $(0 < \lambda \le 1)$, which creates a safety margin. The protection system disconnects the load once the current passes through the feeder exceeds this maximum current; this results in a power outage in the area supplied by this feeder.
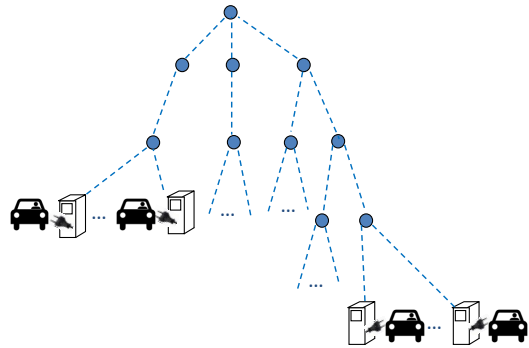

Figure 3: The measurement and signalling architecture.

congestion control schemes). In principle, any one of these schemes can be adopted for EV charger control. Instead of picking one particular scheme, we outline three distributed congestion control schemes that illustrate different points in the design space. Our chosen algorithms differ in the entities that makes decisions about charging rates of EVs and the degree of communication overhead.

3.5.1 Intelligent Endpoint Approach

In this approach, each EV charger independently decides on its charging rate, much like a TCP endpoint. Decision making is *distributed* in this strategy because every EV charger sets its charging rate without having direct knowledge about charging rates of other chargers.

MCC node actions: Every MCC node continuously measures and checks if its corresponding feeder/transfromer is congested (or pre-congested when $\lambda < 1$). Every t_M ms, the root MCC node broadcasts a packet that contains a congestion flag indicating if there is a congestion. This packet is routed hop-by-hop by intermediate MCC nodes until it reaches the EV chargers. Each congested intermediate MCC node can modify the packet that it receives from its parent by setting the congestion flag, if it is not set already, before sending it to all of its children. Therefore, the congestion flag of packets received by an EV charger is set if at least one of its parent MCC nodes detects congestion.

EV charger actions: Every EV charger examines the congestion flag upon receiving a packet from its parent. If an EV charger does not receive any packet in an interval of t_M ms, it acts as if it has received a congestion signal; this is a simple fail-safe mechanism.

Rate allocation algorithm: An EV charger decreases its charging rate to some fraction of its current value every t_C ms $(t_C \ge t_M)$, if a congestion flag was set in at least $\lceil t_C / \alpha t_M \rceil$ of the packets received in the past t_C ms. Otherwise, the charger increases its rate linearly. This is known as the additive increase multiplicative decrease (AIMD) algorithm, which has been thoroughly investigated in the context of the Internet [2].

3.5.2 Local Scheduling Approach

In this approach, the EV charger is slaved to its parent MCC, which makes local scheduling decisions on behalf of the EVs attached to it. Decision making is *distributed* in this strategy similar to the previous scheme; however, it is done by leaf MCC nodes instead of EV chargers. This permits

leaf MCC nodes to discriminate amongst and schedule their downstream EV chargers.

MCC node actions: As before, every MCC node continuously measures and checks if its corresponding feeder or transformer is congested (or pre-congested when $\lambda < 1$). Every t_M ms, the root MCC node broadcasts a packet that contains a congestion flag indicating if there is a congestion. This packet is routed hop-by-hop by intermediate MCC nodes until it reaches the leaf MCC nodes. Each congested intermediate MCC node can modify the packet that it receives from its parent by setting the congestion flag, if it is not set already, before sending it to all of its children. A leaf MCC node infers congestion after receiving a packet if the congestion flag of this packet is set or its own local measurements imply that its corresponding feeder/transformer is congested.

EV charger actions: Every EV charger periodically sends its charging rate and additional information, such as its charging deadline, to its parent MCC node. When it receives a control message from its parent MCC node it sets its charging rate to the rate allocated by its parent and encapsulated in this packet. A fail-safe mechanism requires an EV charger who has not received any packet from its parent in an interval of t_C ms to start backing off its charging rate.

Rate allocation algorithm: A leaf MCC node classifies its downstream EV chargers into a few classes based on their service requirements (for instance, some EVs may have higher priority than others). It allocates its currently available charging rate (which can be determined using an AIMD-style algorithm) amongst its child EVs using a local scheduling approach. For example, if the currently available charging rate was 200A, it could allocate two EVs 80A each, and the remaining two EVs 20A each. If, due to the receipt of a no-congestion signal, the rate were to increase to 200*1.1 = 220A, the MCC could increase the allocated rates to 90A, 90A, 20A, 20A respectively. This novel local scheduling approach combines the benefits of explicit rate allocation with distributed implicit rate computation.

3.5.3 Distributed Explicit Rate Control

In this approach, all MCCs coordinate to select a charging rate for their subtree, in an attempt to minimize oscillations. This approach loosely draws on XCP congestion control [10] Decision making is done *hierarchically* in this strategy because each MCC node can reduce the charging rate of its descendent EV chargers to avoid congestion.

EV charger actions: Every EV charger sends a packet toward the root every t_C ms to negotiate its charging rate for this control interval. This packet contains the current charging rate along with the requested next charging rate of the charger, which is initialized to its maximum charging rate. When this packet returns to the EV charger, it adjusts its charging rate to the charging rate encapsulated in the packet.

MCC node actions: When an MCC node receives a rate requested packet of an EV charger, it may reduce the request rate if its corresponding feeder is congested and this rate is higher than the max-min fair share of this EV charger. Then, it forwards all packets that it has received to its parent. When a rate request packet arrives at the root MCC node, the root sends it back to the EV charger along the same path.

Rate allocation algorithm: Every MCC node waits for a certain amount of time to receive rate request packets from its descendent EV chargers. It then checks if this rate allocation congests its corresponding feeder by comparing the difference between the current available capacity and the sum of current charging rates of its descendent EV chargers with the sum of negotiated rates of these EV chargers. If this rate allocation congests the feeder, the MCC node reduces the negotiated rates of some EV chargers using, for example, the max-min fair share allocation algorithm. Therefore, the most congested feeder along the path always determines the charging rates of corresponding EVs. Note that the negotiated charging rate of an EV charger never increases when the rate request packet passes intermediate MCC nodes to reach the root.

3.5.4 Comparison

We now compare the three schemes outlined in this section.

- *Who has control?*: In the first scheme, control is distributed among EV chargers. However, the second and the third schemes cede control to the utility because rate allocation is done by MCC nodes that are owned and protected by it. Policing the charging rates is required when EV chargers make decisions individually as in the first scheme (since a natural assumption is, in that case, that the chargers are not owned by the utility) since non-compliant chargers may compromise the reliability of the grid. This policing may be as simple as the installation of a circuit breaker that limits the charging current to some pre-defined maximum. If EV chargers are owned by the utility this is not necessary as long as the chargers are tamper-resistant.
- *Oscillations*: The third scheme tries to minimize oscillations by accurate and continuous computation of the remaining capacity and doing rate allocation on this basis. Studies have shown that, at least in the Internet, this reduces the degree of oscillations compared to a standard AIMD approach.
- *Communication Overhead*: The third scheme has a higher communication overhead because control packets travel bi-directionally rather than unidirectionally, as they do in the first two schemes.

It is evident that the schemes differ in many critical aspects: this is the reason why we chose them to illustrate different points in the design space. Yet, all three schemes conform to the basic paradigm of fast, distributed control of EV charging.

4. DISCUSSION AND CONCLUSION

We have outlined the use of real-time distributed control algorithms, based on congestion control algorithms used in the Internet, for the control of charging rates of electric vehicles. We find that the use of low-latency communication allows us to move away from prediction-based approaches and centralized control.

Despite the promising direction of this work, we are aware of two clear challenges that we have not yet addressed. First, transmitting control packets every few milliseconds poses a heavy control overhead. In the Internet, this overhead is amortized by re-using fields in packet headers. In the grid, however, there is no such ongoing communication. Therefore, in future work, we plan to investigate techniques to reduce this onerous communication overhead.

Second, we realize that it is necessary to characterize the effectiveness of the three "control knobs," namely λ, t_M, and t_C on system stability, fairness, and robustness. Their optimal values depend on the critical timescale of our system, which is the timescale of the protection system (*i.e.*, how fast protective relays trip once congestion is detected). For example, suppose that the communication delay in a local area is denoted by d, where d is approximately 1 ms. The control system does not overshoot the control target as long as $t_M + t_C + d$ is less than the critical timescale of the protection system. This means that we have the chance to correct an error made by the last control action if we measure the grid status and send it to the controllers before the protective relay trips. However, if $t_M + t_C + d$ is greater than the timescale of the protection system we should set $\lambda < 1$ to provide a safety margin. To choose the value of λ we must obtain an upper bound on the variation in the demand of uncontrolled loads over a period of t_M ms. This is an exciting direction for future work.

5. REFERENCES

[1] Optimal decentralized charging control algorithm for electrified vehicles connected to smart grid. *Journal of Power Sources*, 196(2):10369 – 10379, 2011.

[2] D. Bansal and H. Balakrishnan. Binomial congestion control algorithms. In *INFOCOM 2001. Twentieth Annual Joint Conference of the IEEE Computer and Communications Societies. Proceedings. IEEE*, volume 2, pages 631 –640 vol.2, 2001.

[3] 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, jun 2011.

[4] R. Brown. Impact of smart grid on distribution system design. In *Power and Energy Society General Meeting - Conversion and Delivery of Electrical Energy in the 21st Century, 2008 IEEE*, pages 1 –4, july 2008.

[5] W. Eddy and M. Welzl. Congestion Control in the RFC Series. http://tools.ietf.org/html/rfc5783, 2010.

[6] EPRI. Impact of Plug-in Hybrid Electric Vehicles on Utility Distribution . Technical Report, 2009.

[7] 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 *e-Energy '12*. ACM, 2012.

[8] X. Gong, T. Lin, and B. Su. Survey on the impact of electric vehicles on power distribution grid. In *Power Engineering and Automation Conference (PEAM), 2011 IEEE*, volume 2, pages 553 –557, sept. 2011.

[9] J. Jaffe. Bottleneck flow control. *Communications, IEEE Transactions on*, 29(7):954 – 962, jul 1981.

[10] D. Katabi, M. Handley, and C. Rohrs. Congestion control for high bandwidth-delay product networks. SIGCOMM '02, pages 89–102. ACM, 2002.

[11] J. Wang, D. X. Wei, J.-Y. Choi, and S. H. Low. Modelling and Stability of Fast TCP. In P. Agrawal, P. J. Fleming, L. Zhang, D. M. Andrews, and G. Yin, editors, *Wireless Communications*, volume 143 of *The IMA Volumes in Mathematics and its Applications*, pages 331–356. Springer New York, 2007.