By Zachary J. Lee, Sunash Sharma, and Steven H. Low

Research Tools for Smart Electric Vehicle Charging

An introduction to the adaptive charging network research portal.


©SHUTTERSTOCK.COM/BUFFALOBOY

Millions of electric vehicles (EVs) will enter service in the next decade, generating gigawatthours of additional energy demand. Charging these EVs cleanly, affordably, and without excessive stress on the grid will require advances in charging system design, hardware, monitoring, and control. Collectively, we refer to these advances as *smart charging*.

While researchers have explored smart charging for more than a decade, very few smart-charging systems have been deployed in practice, leaving a sizeable gap between the research literature and real world. In particular, we find that research is often based on simplified theoretical models. These simple models make analysis tractable but do not account for the complexities of physical systems. Moreover, researchers often lack the data needed to evaluate the performance of their algorithms on real workloads or apply techniques like machine learning. Even when promising algorithms are developed, they are rarely deployed since field tests can be costly and time-consuming.

To address these gaps and accelerate research in smart-charging systems, we have developed the Adaptive Charging Network (ACN) Research Portal, which is made up of three parts: a public data set of EV charging sessions; an open source, data-driven simulation environment; and a platform for field-testing algorithms on real charging systems. In this article, we introduce these tools and the research they have enabled.

ACNs

To enable smart charging at scale, we first needed a platform that would allow us to take measurements from the

Digital Object Identifier 10.1109/MELE.2021.3093597
Date of current version: 2 September 2021

physical system, compute control actions, and apply those actions back to the physical system. When we began our work on smart EV charging, these systems did not exist for the scale of smart charging we envisioned, so we decided to build our own, which we called the ACN. Doing so provided us the platform we needed to run experiments and gather data. It also yielded insights into the practical challenges of smart EV charging systems.

ACN Architecture

The ACN is best understood in the framework of cyber-physical systems (Figure 1). Within the physical system, the local electrical infrastructure supplies power for EVs and other loads from a grid connection or local generation (Figure 2). Sensors within the local infrastructure measure voltages, currents, and power and feed those measurements into the information system through the communications interface, which is part of the information system. The ACN uses several communication protocols, including a Zigbee mesh network for communicating with charging stations and TCP/IP over Ethernet to communicate with power meters.

Data storage and control are distributed between an onsite industrial computer and the cloud. The onsite computer provides robustness to Internet outages, while the cloud provides long-term data storage and visualization. Algorithms for prediction, pricing, and control can be run locally or in the cloud. However, we generally prefer local control algorithms to ensure reliability.

Once the system has calculated a new control action, it sends those set points to the EV supply equipment (EVSE). The EVSE is more commonly known as the *charging port* or *charger*, though the latter is a misnomer since the vehicle's onboard charger is used in level 2 ac charging. This EVSE communicates to the vehicle's onboard battery management system (BMS) using a pilot signal that sets an upper bound on the vehicle's current draw from the EVSE. This pilot signal is defined by the J1772 standard.

Because drivers play an essential role within a charging system, determining when their vehicle is available to charge and how much energy it will need, we provide a mobile app to facilitate communication between the system and driver. Using this app, the driver can enter the vehicle's battery size and maximum charging rate. We also ask the user to provide estimates of the departure time and amount of energy he or she would like delivered to the vehicle. The latter is important since the J1772 standard does not

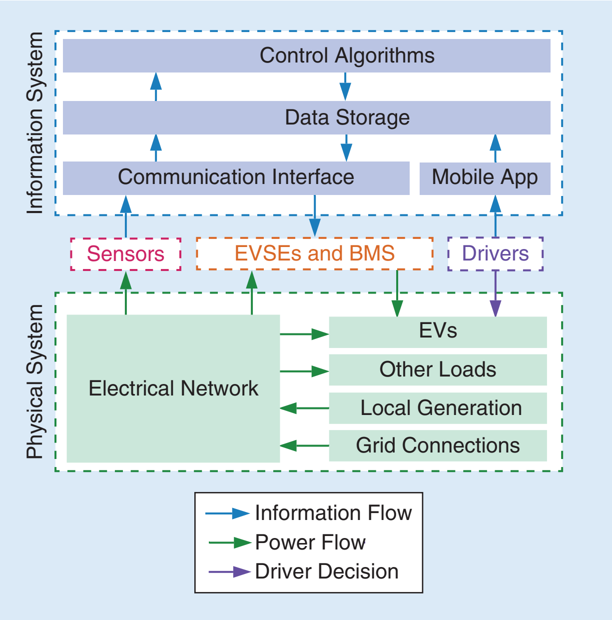

Figure 1. The architecture of ACNs. BMS: battery management system; EVSE: EV supply equipment.

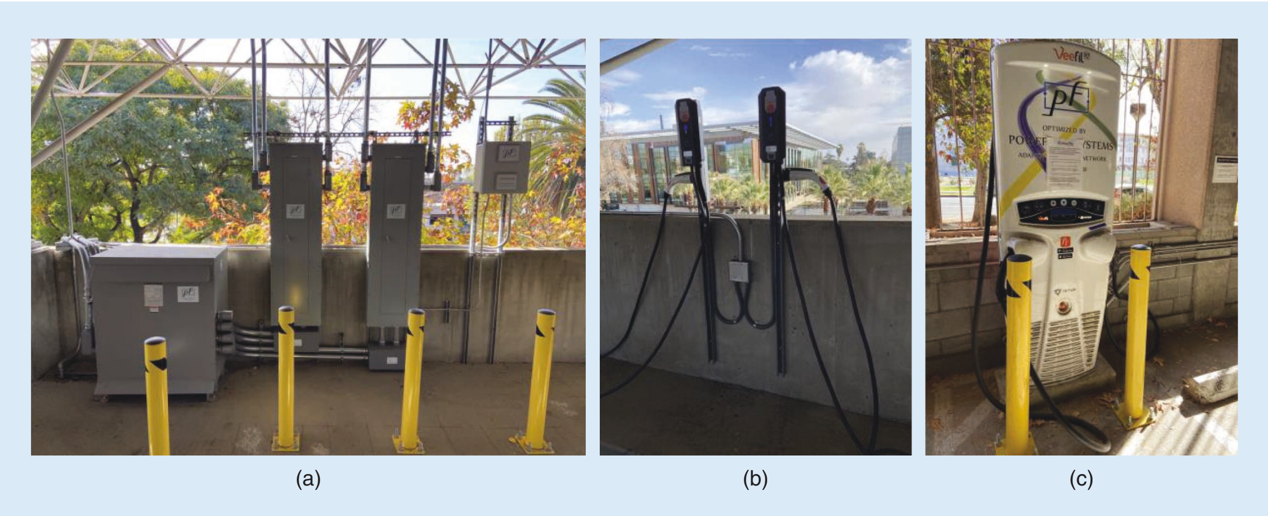

Figure 2. The ACN equipment: (a) from left to right, the constrained transformer, electrical panels, and PowerFlex controller; (b) a Webasto DX level 2 EV supply equipment; and (c) a Tritium Veefil RT50 dc fast charger.

allow for bidirectional communication, so our ACN cannot directly access the vehicle's state of charge for level 2 charging.

Lessons Learned From the ACN

The first ACN was built at the California Institute of Technology (Caltech) in early 2016. Since then, PowerFlex, a company formed to commercialize the technology, has deployed ACNs across the United States. As of May 2021, PowerFlex has deployed more than 4,600 EVSEs at more than 200 sites. By operating these systems over the last five years, we have learned important lessons about large-scale charging systems.

Unbalanced Infrastructure

Within the United States, commercial customers like schools, hospitals, offices, and stores receive three-phase power. Level 2 EVSEs are single-phase loads, so they are connected line to line at these sites. Because of differences in usage at each EVSE, large-scale charging systems can be very unbalanced.

Traditionally, most scheduling and control approaches for EV charging that consider infrastructure limits implicitly assume single-phase or balanced operations, so they only need to consider the power limits of the equipment. However, because of unbalanced three-phase currents, these simple constraints can be insufficient to ensure safe operations.

For example, in Figure 3, we show that a model predictive control (MPC) algorithm that assumes balanced operations can remain below a transformer's aggregate power limit while simultaneously overloading individual lines. However, if we introduce constraints that explicitly model unbalance, we ensure that all line limits are respected.

Infrastructure constraints are expressed as upper bounds on the magnitudes of currents within the system. By Kirchhoff's current law, any line current can be written as a linear combination of load currents in phasor form. We then require that the magnitude of any line current is less than its limit. Mathematically, this can be expressed as a second-order cone constraint, which is convex and can be incorporated into many control algorithms, such as MPC.

Pilot Signals and Battery Behavior

We use the J1772 pilot signal to control the charging rate of each vehicle within the ACN. However, this pilot signal is only an upper bound on the current drawn by the vehicle's BMS. Because the pilot signal is an upper bound, if the pilot sent to the BMS satisfies infrastructure constraints, then the actual current draw will as well. This means that the pilot can ensure safety constraints in oversubscribed systems.

However, because the pilot is only an upper bound, the BMS will sometimes underutilize its allocated pilot signal. Sometimes, this is because the pilot is above the maximum charging rate of the vehicle. Other times, this can be a temporary limit.

For example, we find that many BMSs use a constant-current, constant-voltage charging scheme. Within this scheme, the vehicle can usually accept its full charging rate until it reaches about the 80% state of charge. After this, its maximum charging rate decreases roughly linearly with its state of charge. Charging systems that do not account for this BMS behavior can underutilize their infrastructure capacity and underestimate the time needed to finish charging a vehicle.

Discrete Set Points

EVSEs also impose limits on the pilot signals that they support. For example, the J1772 standard does not allow pilot signals below 6 A (except zero). Also, most commercially available EVSEs support only a discrete set of pilot signals. The granularity of this control can vary widely. For instance, some EVSEs support only 16-A increments, while others have 1-A or even 0.1-A increments. Even at 1-A increments, a simple rounding scheme can leave a significant amount of capacity unused. A site with 50 EVSEs might result in up to 10 kW of wasted capacity due to rounding.

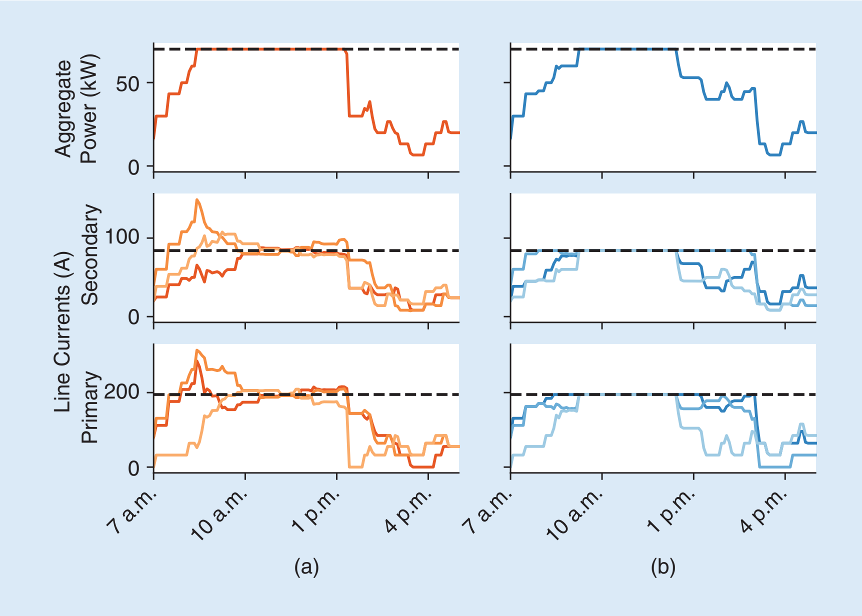

Figure 3. (a) A balanced model is insufficient in large-scale charging systems where unbalance can be significant. (b) Instead, second-order cone constraints that account for unbalance are necessary.

ACN Research Portal

Beyond serving as a model for smart-charging systems, the ACN has led to the creation of the ACN Research Portal (Figure 4). This portal has three parts:

- *ACN-Data*: a collection of fine-grained charging data collected from the Caltech ACN and similar sites
- *ACN-Sim*: an open source simulator that uses ACN-Data and realistic models derived from actual ACNs to provide researchers with an environment to evaluate their algorithms and test assumptions
- *ACN-Live*: a framework for safely field-testing algorithms directly on the Caltech ACN.

ACN-Data

The ACNs at Caltech, NASA's Jet Propulsion Laboratory (JPL), and elsewhere have produced a vast amount of data that have enabled new research lines within our lab. However, most researchers do not have access to data from systems like the ACN, precluding them from applying data-driven methods, such as machine learning and trace-driven simulations, in their work. Instead, most existing works focus on assumed data distributions or data collected from internal combustion engine vehicles.

To meet this need in the community, we have collected and published ACN-Data, which includes data from 207 level 2 EVSEs and six dc fast chargers (DCFCs) from seven clusters: five at Caltech (including one at a Laser Interferometer Gravitational-Wave Observatory facility in Louisiana), one at JPL, and one at an office building in northern California. These clusters cover common use cases, including campus/public-use and access-controlled workplace charging. The smallest of these clusters has only four level 2 chargers, while the largest has 78 level 2 chargers and two DCFCs in a single parking structure.

The data set includes more than 80,000 sessions. Each session includes the arrival time, departure time, energy delivered, and user data from a mobile app, including the vehicle information, estimated departure time, and estimated energy request. The data set also includes time series data, including the charging current, power, and voltage, with 4-s resolution.

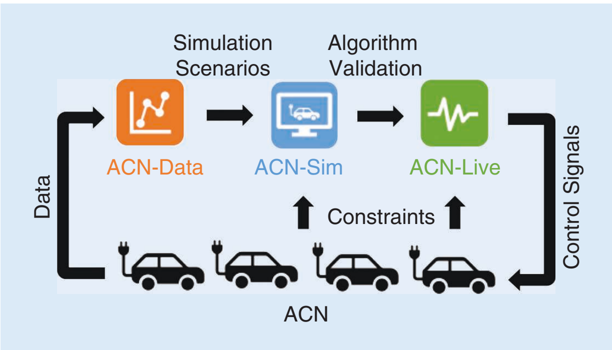

Figure 4. An overview of the ACN Research Portal.

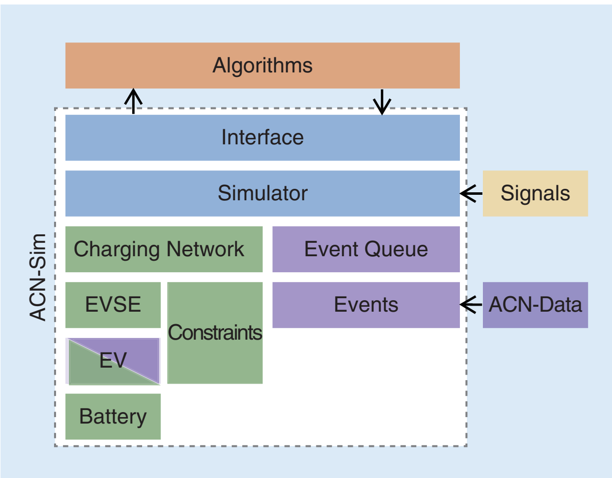

Figure 5. The modular architecture of ACN-Sim.

ACN-Sim

While the ACN has allowed us to identify many interesting challenges in real-world EV charging systems, we recognize that most researchers will not have access to such physical testbeds. However, it is still important that they have a realistic environment in which to evaluate their algorithms. To this end, we developed ACN-Sim, a modular, data-driven simulation environment for testing scheduling algorithms for EV charging systems (Figure 5).

- *Models*: ACN-Sim includes realistic models of the many components of a real EV charging system, such as electrical infrastructure, charging stations, and EV BMSs. These models incorporate our lessons from real systems. For example, the charging network constraints include three-phase unbalanced models, while batteries use the constant-current, constant-voltage charging scheme discussed previously. The simulator is also designed to be highly modular, meaning that users can replace each component to model different types of hardware or levels of fidelity.
- *Events*: The simulator is event driven. To generate events, users can get real event sequences from ACN-Data, generate them from statistical models, or manually create events to investigate edge cases. To make accessing ACN-Data simpler for users, ACN-Sim provides direct integration with the ACN-Data application programming interface (API). This allows the user to specify a site and date range, and ACN-Sim will gather the actual workload from that ACN and generate the appropriate plug-in and unplug events. ACN-Sim also provides utilities for learning statistical models, such as Gaussian mixture models (GMMs), directly from data using tools from scikit-learn.
- *Signals*: The signals submodule allows ACN-Sim to integrate with external signal sources, which can be an important part of EV charging systems such as utility tariffs, solar-generation curves, and external loads. These signals are available to the scheduling algorithm through the simulator's interface.

- *Interface*: To make algorithm implementations more flexible, we introduce an interface that abstracts away the underlying infrastructure, whether simulated or real, allowing us to use the same algorithm implementation with both ACN-Sim and ACN-Live. This means that users can thoroughly test algorithms with ACN-Sim before trying on physical hardware.
- *Defining algorithms*: To define an algorithm in ACN-Sim, users only need to extend the BaseAlgorithm class and define the *schedule()* function. This function takes in a list of active sessions—meaning that the EV is plugged in, and its energy demand has not been met—and returns a charging schedule for each. To make baselining new algorithms easier, ACN-Sim is packaged with common scheduling algorithms, such as round robin (RR) (equal sharing); first come, first served (FCFS); earliest deadline first (EDF); and least laxity first (LLF). We also provide the adacharge package, which includes a flexible MPC framework designed to work with ACN-Sim.
- *Integrations*: ACN-Sim is designed to integrate with other packages, including the OpenAI gym for reinforcement learning and grid simulators, such as OpenDSS, PandaPower, and MATPOWER.

ACN-Live

Thus far, we have discussed how ACN-Data and ACN-Sim help researchers apply data-driven methods to EV charging research. While these tools help bridge the gap between theoretical work and algorithms that could be deployed in practice, they are not substitutes for field tests and pilots. These field tests are important to prove to stakeholders, like utilities, funding agencies, policy makers, and consumers, the viability of new technologies.

However, field tests and pilots in the energy space are rare, as they require vast amounts of time, funding, and expertise. For example, the ACN system has required more than five years of work and millions of dollars of funding to reach its current state. Because of these challenges, only a handful of researchers have access to systems like the ACN.

This lack of access hampers research progress and technology transfer into the marketplace. To bridge this gap, we have designed ACN-Live, a framework for field-testing algorithms on the Caltech ACN. ACN-Live allows researchers who have thoroughly tested their algorithms with ACN-Sim to deploy them on the physical ACN. By utilizing the same interface as ACN-Sim, we enable researchers to perform field tests with no changes to their algorithm implementation.

This is a unique opportunity that requires the specialized hardware of the ACN and close collaboration among our research group, the Caltech facilities, and PowerFlex. As such, we are likely the only facility in the world able to provide this type of hardware-in-the-loop testing to the research community.

Adaptive Scheduling Algorithms

These tools have allowed us to develop practical algorithms for large-scale, smart EV charging. Though there is vast literature on charging algorithms, we have found that most make strong assumptions, which prevents us from using them in practice. Others lack the flexibility to incorporate practical constraints and objectives.

Algorithm Design

To account for these practical needs, we developed an algorithm based on MPC. We refer to this algorithm as the *adaptive scheduling algorithm* (ASA). At each time step, we form an optimization problem. The objective of this optimization can incorporate many operator objectives, such as minimizing the cost, charging quickly, or flattening the load. We also introduce a collection of regularizers that promote desirable properties in the final schedule, such as fairness or smoothing.

The constraints of the optimization express limits on the state and action spaces of the problem. For example, we introduce limits on the charging rate of the vehicle, unbalanced three-phase infrastructure constraints, and constraints on the energy delivered to each vehicle.

Accounting for Discrete Pilots

We do not directly include the discrete pilot signals as constraints in this optimization. Doing so would make the problem a discrete optimization, which may be intractable to solve within the time constraints of MPC. Instead, we require that the charging rate of each vehicle lie between zero and the upper limit of the EVSE. We then perform a rounding and capacity reallocation step in postprocessing.

Ensuring Feasibility

We also do not require that the energy delivered to each vehicle match the energy requested by the user, which might lead to infeasibility in the optimization problem. Instead, we require that the energy delivered to each vehicle is less than the amount requested by the user. We then add a penalty term to the objective, ensuring that this inequality is tight when fully meeting the energy demand is feasible.

Reclaiming Idle Capacity

As we saw previously, an EV's BMS will sometimes limit the power draw of the battery as it approaches a 100% state of charge. When this happens, the difference between the pilot signal and vehicle's actual charging rate is wasted capacity. To reclaim this capacity, we use a simple algorithm that we call *ramp down*.

When the charge rate of the vehicle is less than a given threshold below the pilot, we decrement the upper limit on the pilot signal to reclaim some capacity. Likewise, when the charging rate is near the pilot, we increment this upper bound. With this scheme, we can quickly reclaim capacity as the battery fills up, while still allowing EVs to throttle back up if this reclamation was premature.

Importance of Unbalanced Three-Phase Models

Unbalance can be a significant concern in large-scale charging systems. However, to date, most algorithms proposed in the literature implicitly assume single- or balanced three-phase operation. As we have seen, unbalanced three-phase constraints are necessary to ensure safety. These unbalanced models can significantly impact the performance of an algorithm. To see this, we can use the ACN Research Portal to evaluate the percentage of user energy demands met when using balanced and unbalanced models.

We compare six algorithms over a range of possible transformer capacities based on the actual charging workload of the ACN at Caltech in September 2018. For this experiment, we use our ASA with an objective that promotes charging as quickly as possible, which we denote ASA-QC.

From Figure 6, we can see that in the balanced case, EDF, LLF, and MPC all perform near optimally, exceeding the performance of RR and FCFS by up to 8.6%. However, in the unbalanced case, we see that, while ASA-QC can match the offline optimal as before, EDF and LLF both underperform. In the highly constrained regime, RR outperforms EDF and LLF despite having less information about the workload. We attribute these results to the importance of phase balancing in three-phase systems, which has been historically underappreciated in the managed charging literature.

To understand why ASA performs so much better than the baselines, we must consider what information each algorithm uses. RR uses how many EVs are present and performs the worst. EDF uses only information about departure time, while LLF also uses the EV's energy demand. Only ASA-QC actively optimizes over infrastructure constraints, allowing it to better balance phases (increasing throughput) and prioritize EVs, including the current and anticipated congestion. A key feature of the ASA framework is its ability to account for all available information cleanly.

Figure 6 can also be used to evaluate the infrastructure needs of a site. For example, we can see that, if a host wants to deliver >99% of charging demand using ASA-QC, a 70-kW transformer would be sufficient, assuming an unbalanced three-phase system. Alternatively, if an existing transformer can only support 40 kW of additional demand, a host could expect to meet approximately 85% of energy demands without an upgrade.

Interfacing With the Grid

Charging systems do not operate in a vacuum. In almost all cases, they draw energy from the power grid. Because of their enormous power and energy requirements, large-scale EV charging systems can significantly impact the power grid.

The ACN Research Portal allows us to expand studies like this to consider how more advanced smart-charging approaches can help alleviate strain on the distribution system, especially for large charging systems like workplaces. To enable studies like this with ACN-Sim, we have integrated it with several grid simulation packages, including MATPOWER, PandaPower, and OpenDSS. In each case, we can use ACN-Sim and ACN-Data to provide a realistic load profile, which can then feed into the grid simulation package that evaluates power flows and alerts us to any voltage or overloading issues.

For this case study, we use OpenDSS to model a 240-node distribution system in the midwestern United States with hourly smart meter data. Our goal is to evaluate the effect of installing a 52-EVSE charging network at one node in this distribution system. We use actual data from the ACN at JPL on 6 September 2018. There are four cases: a baseline with no EV charging, uncontrolled charging, ASA with a load-flattening objective, and ASA with load flattening and onsite solar.

The results of this experiment are shown in Figure 7. We can see that uncontrolled charging at this scale would overload the distribution transformer and lead to unacceptably low voltages in the network. However, using ASA with load flattening, we can stay below the transformer capacity and above the voltage limits. Moreover, we see that if 225 kW_{ac} of solar were installed at the site, we achieve the same circuit-wide minimum voltage as the baseline case. This indicates that onsite solar generation and smart EV charging could enable widespread workplace EV charging without adverse grid impacts.

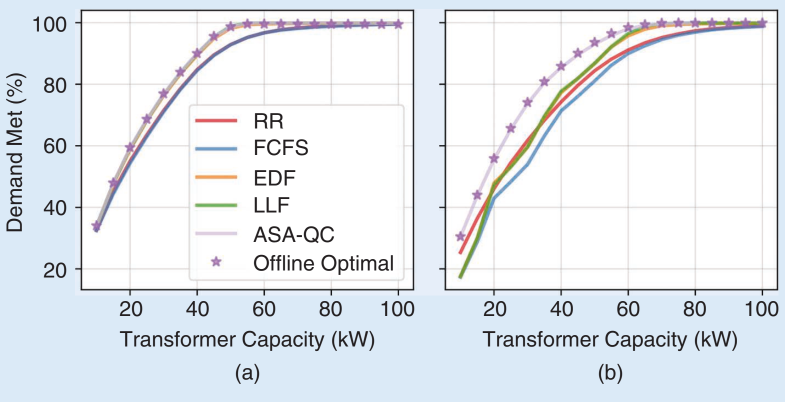

Figure 6. The algorithm performance with constrained infrastructure with the (a) balanced and (b) unbalanced models.

Data-Driven Modeling

In addition to enabling trace-driven simulation, we can also use the ACN Research Portal to develop statistical models. These models can help us predict user behavior and evaluate charging system designs before they are built.

Modeling Workloads With GMMs

There are many approaches to modeling charging workloads, including kernel density estimation and normalizing flow methods. In our case, we consider GMMs to jointly model the arrival time, session duration, and energy delivered in each charging session. In this context, the GMM lends itself to a natural interpretation.

We assume that most EV drivers have a finite set of normal routines. For example, a driver might typically plan to arrive at work each morning at 8 a.m. and remain for 9 h, leaving at 5 p.m. This driver also likely follows a similar route to work, so his or her energy needs are similar day to day. However, this driver's actual arrival time, duration, and energy request are influenced by noise, such as traffic. It is also natural to assume that these variables are correlated; i.e., more traffic might mean a later arrival time and higher energy needs. These intuitions motivate us to consider modeling the driver's routines as a multivariate Gaussian distribution.

Drivers may also have several routines. For example, a driver may drop his or her kids off at school on some mornings but not others. We can model this by considering a mixture of multivariate Gaussians. We can assume that drivers may share similar routines, i.e., commuting to work from similar areas. Therefore, we also consider population-level models with far fewer components (Gaussians) than the number of drivers in the population.

Predicting Session Parameters

After fitting a GMM to historical data, session parameters are estimated by conditioning this distribution on the vehicle's arrival time and taking the expectation of the remaining two variables. We can make this prediction using models trained on individual data or at the population level. To account for changing behavior over time, the GMMs are periodically refit using data from a rolling window. Experiments show that a window of 30–60 days provides a good tradeoff between data quantity and quality.

Even these simple models provide better accuracy than user inputs through the mobile app, achieving a symmetric mean absolute percentage error of less than 12.3% for duration and 12.8% for energy. Meanwhile, user inputs have errors of 18.6% for duration and 26.9% for energy. However, there is still considerable work to be done to increase the prediction accuracy and account for other covariates.

Evaluating Charging System Designs With Data

Smart charging can significantly reduce the capital and operating costs of charging systems. We can use ACN-Data and ACN-Sim to quantify these benefits for a particular use case.

Consider a site host who wants to install an EV charging solution at an office building. The host estimates that the system will charge approximately 100 EVs per day. Several potential designs can be considered, as shown in Table 1.

We assume that the office will have a usage pattern similar to that of JPL, so we train a GMM based on the data collected on weekdays at JPL. We assume the site will not allow usage on weekends. We then use ACN-Sim's *GaussianMixtureEvents* tool to create a queue of events from this generative model, assuming 100 arrivals on weekdays and zero on weekends. Since EVs are generated, we use ACN-Sim's *StocasticNetwork*, which randomly assigns EVs to EVSEs when they arrive, to model each design.

To evaluate costs, we use the Southern California Edison EV TOU-4 tariff, which includes a time-of-use energy tariff and demand charge on peak power draw. The experiments are repeated for 10 months of generated data, with the mean results shown in Table 1. Note that the standard deviations among months were less than 3.5% for each metric in each case.

From Table 1, we see that, while installing 100 level 1 EVSEs might be the simplest solution, these slow chargers can meet only 75.4% of demand because they cannot support users with large energy needs and short deadlines. However, the alternative of installing a

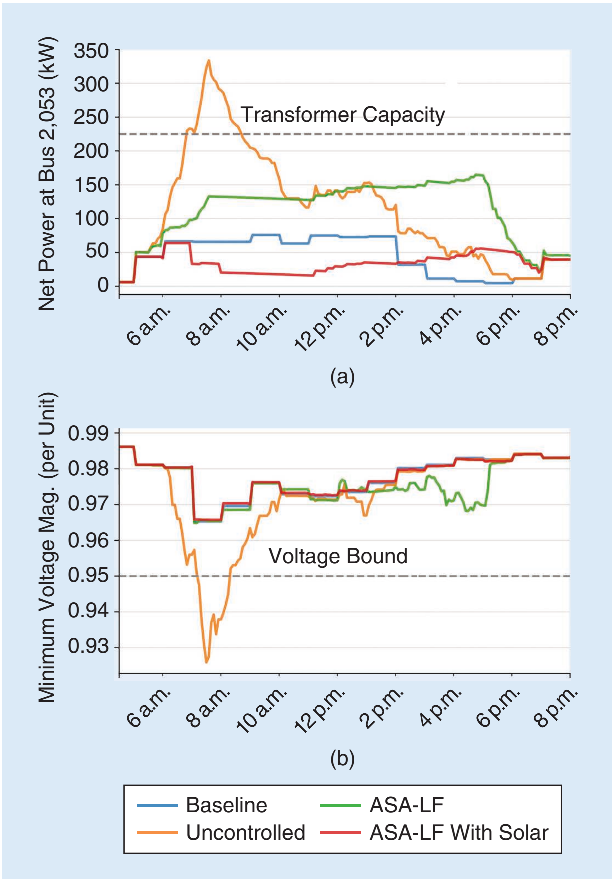

Figure 7. Comparing the effect of (a) controlled and (b) uncontrolled charging on net power at the bus and system-wide minimum voltage. LF: load flattening; Mag.: magnitude.

TABLE 1. The infrastructure solution evaluation at 100 EV/day.

EVSEs	EVSE Level	Algorithm	Capacity (kW)	Swaps (Number/Month)	Demand Met (%)	Cost (US$/kWh)
102	Level 1	Unctrl	200	0	75.4	0.278
102	Level 2	Unctrl	680	0	99.9	0.351
30	Level 2	Unctrl	200	1,103.5	99.6	0.256
102	Level 2	ASA-CM	200	0	99.8	0.234

Unctrl: uncontrolled; CM: cost minimization.

TABLE 2. The infrastructure solution evaluation at 200 EV/day.

EVSEs	EVSE Level	Algorithm	Capacity (kW)	Swaps (Number/Month)	Demand Met (%)	Cost (US$/kWh)
102	Level 1	Unctrl	200	1,174.5	73.2	0.244
102	Level 2	Unctrl	680	1,081.5	99.8	0.327
30	Level 2	Unctrl	200	2,973.9	91.6	0.233
102	Level 2	ASA-CM	200	1,441.9	87.1	0.223
201	Level 2	ASA-CM	200	0	98.4	0.227

680- kW transformer and the associated service upgrade would be cost prohibitive for most sites, and installing only 30 level 2 EVSEs requires more than 1,100 swaps per month, leading to lost productivity and poor user experience. In this case, we see that ASA with the objective of cost minimization (CM) results in reduced capital costs, higher user satisfaction, and lower operating expenses.

The benefits of smart charging are amplified as EV adoption grows, and charging infrastructure must scale accordingly. In this scenario, we consider how the system will rise to 200 charging sessions per day. The results are shown in Table 2. Intuitively, the systems designed for 100 EVs per day require far more swaps with increased demand, and, similarly, the percentage of demand met decreases. This is also true for the smart charging (ASA-CM) case.

However, while scaling the number of EVSEs in traditional uncontrolled charging systems would require a corresponding scaling of the transformer capacity to ensure safety, smart charging enables us to add new EVSEs without increasing the transformer capacity. This allows us to meet more than 99.8% of energy demands, require zero swaps, and maintain low operating costs.

Conclusion

Smart charging will be key to enabling safe and cost-effective EV charging at scale. However, new research and development will be needed to bring the promises of smart charging to the marketplace. The ACN Research Portal provides researchers the data and tools they need to develop these new technologies and deploy them quickly.

For Further Reading

Z. J. Lee et al., "Adaptive charging networks: A framework for smart electric vehicle charging," *IEEE Trans. Smart Grid*, early access, 2021. doi: 10.1109/TSG.2021.3074437.

Z. J. Lee, S. Sharma, D. Johansson, and S. H. Low, "ACN-sim: An open-source simulator for data-driven electric vehicle charging research," *IEEE Trans. Smart Grid*, early access, 2021. doi: 10.1109/TSG.2021.3103156.

Z. J. Lee, J. Z. Pang, and S. H. Low, "Pricing EV charging service with demand charge," *Electric Power Syst. Res*., vol. 189, p. 106,694, Dec. 2020. doi: 10.1016/j.epsr.2020.106694.

Z. J. Lee, T. Li, and S. H. Low, "ACN-data: Analysis and applications of an open EV charging dataset," in *Proc. 10th ACM Int. Conf. Future Energy Syst.*, 2019, pp. 139–149. doi: 10.1145/3307772.3328313.

Z. J. Lee and S. Sharma, "ACN portal," GitHub, 2021. https://github.com/zach401/acnportal

Z. J. Lee and S. Sharma, "ACN research portal examples," GitHub, 2021. https://github.com/caltech-netlab/acnportal-experiment

Biographies

Zachary J. Lee (zlee@caltech.edu) is with the Department of Electrical Engineering, California Institute of Technology, Pasadena, California, 91125, USA, and PowerFlex Systems, Los Altos, California, 94022, USA.

Sunash Sharma (sbsharma@caltech.edu) is with the Department of Computing and Mathematical Sciences, California Institute of Technology, Pasadena, California, 91125, USA.

Steven H. Low (slow@caltech.edu) is with the Departments of Electrical Engineering and Computing and Mathematical Sciences, California Institute of Technology, Pasadena, California, 91125, USA.