Secure Integration of Electric Vehicles with the Power Grid

Chaitra Niddodi*, Shanny Lin+, Sibin Mohan*, and Hao Zhu+

*University of Illinois at Urbana-Champaign, +The University of Texas at Austin

***Abstract*—This paper focuses on the secure integration of distributed energy resources (DERs), especially pluggable electric vehicles (EVs), with the power grid. We consider the vehicle-to-grid (V2G) system where EVs are connected to the power grid through an 'aggregator' In this paper, we propose a novel *Cyber-Physical Anomaly Detection Engine* that monitors system behavior and detects anomalies almost instantaneously (worst case inspection time for a packet is 0.165 seconds[1]). This detection engine ensures that the critical power grid component (*viz.,*aggregator) remains secure by monitoring *(a)* cyber messages for various state changes and *data constraints* along with *(b)* power data on the V2G cyber network using power measurements from sensors on the physical/power distribution network. Since the V2G system is time-sensitive, the anomaly detection engine also monitors the *timing requirements* of the protocol messages to enhance the safety of the aggregator. To the best of our knowledge, this is the first piece of work that combines *(a)* the EV charging/discharging protocols, the *(b)* cyber network and *(c)* power measurements from physical network to detect intrusions in the EV to power grid system.**

***Index Terms*—Cyber-Physical, Security, Intrusion Detection, Vehicle-to-Grid, Electric Vehicles, Anomaly Detection**

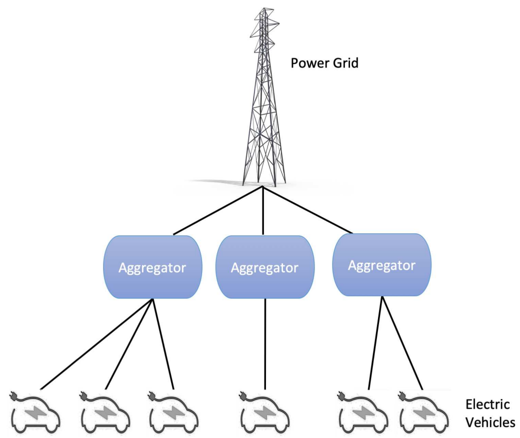

Fig. 1: Conceptual Architecture of a Vehicle-to-Grid (V2G) system. [15]

I. Introduction

Fig.1 shows the conceptual architecture of the V2G system. The main components in this system include EVs, aggregators and the power grid. The EV to power grid operations considered in this paper are *(a)* charging where EV acts as a load and draws power from the grid, *(b)* discharging where EV acts as a power generator and supplies power to the grid. Aggregators are entities that act as mediators between end users (*viz.,*EVs) and the utility operator [7]. Aggregators are particularly useful in coordinating discharging operations between the EVs and the power grid. This is because individual EVs have very small power capacities in comparison with the scales of power generation and distribution at the power grid. Therefore, for efficient discharging operations, a large number of EVs are required. With aggregators acting as intermediaries between the utility power grid operator and the EVs [7], all communication messages between the EVs and the power grid pass through aggregators. In the model presented in Figure 1, the aggregator can be a prime target for attackers since *(i)* it manages multiple EVs and *(ii)* is also directly connected to the power grid utility system. A successful intrusion at the aggregator level can have serious consequences for the power grid. In fact, it is well documented that the power grid is vulnerable to a wide range of attacks [13]. Therefore, ensuring the security of this critical component (*viz.,*aggregator) is essential to ensure secure integration of DERs such as EVs with the power grid. To this end, we *propose a Cyber-Physical Anomaly Detection Engine* with mechanisms to detect anomalous behavior in aggregators of the V2G system. For our anomaly detection engine, we rely on both the cyber and physical properties of the system. On the cyber side, we focus on the communication protocol in the V2G system to ensure correct operation of the aggregator, while we validate its behavior using the physical side of the system in the form of power measurements.

The main **contributions** of this work are:

1) An enumeration of the correct sequences of commands in the V2G communication protocol. This is used to generate an aggregator state machine for our detection engine (§IV-A)
2) Development of a Cyber-Physical Anomaly Detection Engine that can detect malicious activity at the aggregator level as soon as they occur. The anomaly detection engine monitors communication on the V2G cyber network using power measurements from sensors on the physical/power distribution network (§IV-D). It also makes use of *timing constraints* related to frequency of periodic messages (§IV-B) and subscription period (§IV-C) to differentiate between correct/incorrect system behavior.
3) Implementation and evaluation of a prototype of the Cyber-Physical Anomaly Detection Engine (§V)

[1]Minimum latency on V2G network is 2 seconds

978-1-5386-8099-5/19/$31.00 ©2019 IEEE

978-1-5386-8099-5/19/$31.00 ©2019 IEEE

While there are some intrusion detection systems designed for components of the power grid system (as discussed in detail in §VII), to the best of our knowledge there are none that combine cyber and physical properties of the system along with communication standards for EVs. Hence, there is no direct comparison possible while evaluating our cyber-physical anomaly detection engine. Our evaluation consists of measuring the accuracy and performance of our anomaly detection engine as described in §V. The simple model of our anomaly detection engine enables it to detect anomalies accurately and almost instantaneously.

II. System Model

Fig.2 highlights the system under consideration along with the various components and relevant connections in the cyber and physical networks. Multiple buildings (*e.g.,* households) are connected to the aggregator with multiple EVs connected to each building. The Cyber-Physical Anomaly Detection Engine residing at the aggregator receives inputs from both the cyber as well as physical networks. We assume a second cyber network connecting the sensors in the physical network / power distribution network to the Cyber-Physical Anomaly Detection Engine in order to receive power measurements. For the purpose of evaluating the performance of our system, we assume that this second cyber network has similar properties (in terms of bandwidth etc) as the cyber network in the V2G system.

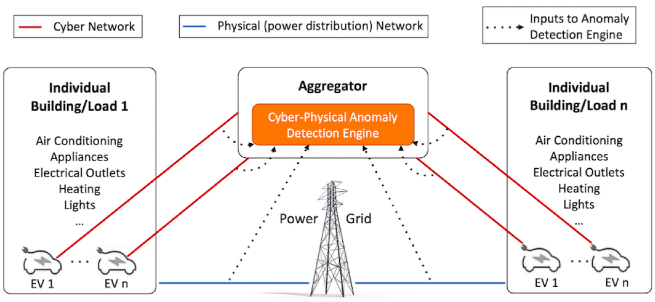

Fig. 2: System Model: EVs intermittently charging at homes

As part of the cyber network, the EVs are the only type of load being considered, i.e, only the packets exchanged between the EVs and aggregator are taken as input for anomaly detection. As part of the physical network, the aggregated profile of all household loads at the entry point is taken as input for anomaly detection. Individual household loads such as air conditioning, appliances, electrical outlets, heating and lighting loads etc. make up the majority of the non-EV loads in the aggregated profile. In this model, we assume that the physical network and cyber network are not compromised simultaneously, *i.e.,* we assume that the inputs obtained from atleast one of these networks are genuine.

III. Threat Model

V2G system is susceptible to a wide range of attacks including compromise of the individual components as well as the communication channel interconnecting them.

Tab. I shows the attacks that our Cyber-Physical Anomaly Detection engine focuses on. There have been instances of attacks involving modification of current to cause a fire [19]. Multiple studies [9], [10] have shown that EVs can offer significant services to improve grid stability. It therefore follows that compromising a large number of EVs and in turn a large number of aggregators using the above attacks (Tab. I) could cause grid instabilities. With reference to the power grid, security issues manifest as safety/reliability concerns where attackers try to bring down system reliability. Hence ensuring the reliable functioning of the component that connects EVs to the power grid *viz.,*aggregator is of paramount importance for the safe operation of the power grid.

TABLE I: Threat Model

No.	Attack	Effect
1	Compromise of EVs, the cyber network or the power distribution network to report more power than actually consumed.	DOS attack at the aggregator preventing more EVs from connecting to aggregator and causing unexpected variations in power frequency at the aggregator.
2	Compromise of EVs, the cyber network or the power distribution network to report less power than actually consumed.	Could cause transformer overheating due to more EVs connecting to aggregator and also cause unexpected variations in power frequency at the aggregator.
3	Compromise of EVs or the cyber network causing packets to be generated out of expected sequence. *E.g.,* Man-in-the-middle attack on the cyber network where fake packets are sent from EV to aggregator requesting for operation cancellation partway through the operation.	Prevents EVs from completing their charging/discharging operations and causes unexpected variations in power frequency at the aggregator.
4	Compromise of EVs or the cyber network causing EVs to charge/discharge beyond their subscription periods. *E.g.,* Man-in-the-middle attack on the cyber network where a packet from aggregator to EV containing the approved subscription period is modified. In response, EV starts charging when it is not supposed to.	Aggregator instability due to excess power demand.
5	Compromise of EVs or the cyber network causing periodic message packets to be generated more or less frequently than expected. *E.g.,* EV sends very few power status update packets (less frequently than expected).	Inaccurate estimation of power profile at the aggregator.

IV. Cyber-Physical Anomaly Detection Engine

The anomaly detection engine includes *(a)* message sequence validation *(b)* message frequency validation *(c)* subscription period validation and *(d)* power measurement validation to detect anomalies in system behavior as shown in Fig.3. An anomaly is detected whenever the system deviates from the expected system behavior. Expected system behavior is defined based on the communication standards in the V2G system, as discussed in detail below.

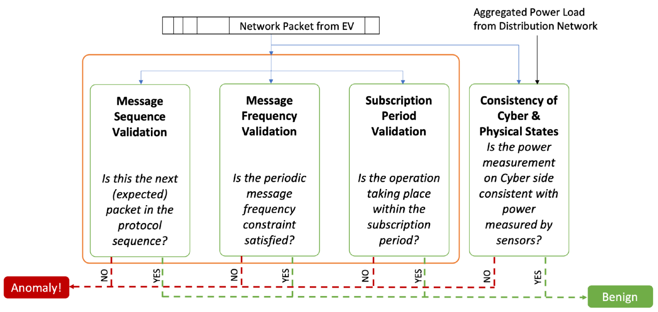

Fig. 3: Components of Cyber-Physical Anomaly Detection Engine

A. Message Sequence Validation

The sequences of messages in incoming packets are validated. To do this, an aggregator state machine is created with valid states and state transitions. This state machine is based on valid message sequences established by communication standards between the aggregator and the EVs. The requirements and specifications for communication between EVs and the electric power grid are established by the SAE J2847/1 standard for forward power flow that includes charging [8] and SAE J2847/3 standard for reverse power flow that includes discharging [8].

Since the SAE standards are proprietary, complete details are not provided in this paper. However, sufficient details on the types and sequences of messages are provided in Fig. 4 for a better understanding of our work.

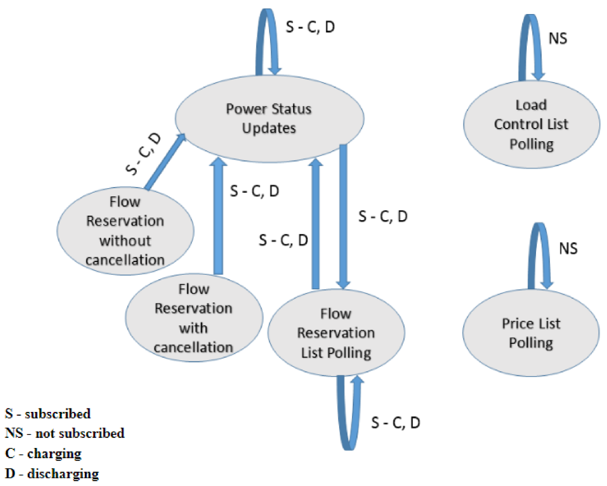

Fig. 4: Sub-states of the aggregator state machine.

Fig. 4 shows the various messages in each of the two EV-grid operations, charging and discharging. It also shows the principal sub-states in the aggregator state machine that are generated from the list of valid command sequences enumerated for the EV-grid operations. Each state represents a *sequence* of messages. The state diagram captures the following states and state transitions. *(a)* EVs are subscribed *i.e.,* connected to the power grid and drawing or supplying power (denoted by 'S'). Flow reservation is the process in which the EV is assigned a subscription period (*i.e.,* period when connected to the grid) for charging/discharging. First, a flow reservation is established. This may be followed by one or more new flow reservations after cancellation of a previously established flow reservation. Once the subscription period begins, the EV starts sending periodic power status updates (*i.e.,* information related to the amount of power drawn) to the grid through the aggregator. In parallel, it also periodically fetches the flow reservation list from the grid, through the aggregator, to check for any updates in the subscription period. *(b)* EVs are not subscribed *i.e.,* connected to the power grid but not drawing or supplying power (denoted by 'NS'). During this phase, EVs only periodically fetch updates on pricing and load control related information from the grid to make a decision on when to charge / discharge. This component of the anomaly detection engine handles attack 3 in Tab. I.

B. Message Frequency Validation

Power grid systems such as V2G system must satisfy certain time related constraints. Our anomaly detection engine monitors incoming packets to ensure that timing constraints are enforced. Periodic messages have a predefined frequency. The frequency of such periodic messages coming into the aggregator are monitored. At the aggregator, the frequency of a given periodic message is monitored by checking the time elapsed between two occurrences of the message. Therefore, it is not required that the clocks at the aggregator and the EVs be synchronized. This component of the anomaly detection engine handles attack 5 in Tab. I.

C. Subscription Period Validation

Message data in the packets coming into the aggregator are validated. Analysis of message data in packets is particularly useful for monitoring the aggregator/EV communication that involves a highly vulnerable edge device *viz.,*the EV. Two important parameters in the EV to power grid communication that are most likely to be tampered by adversaries are:

a) Subscription Period, that defines the duration of charging/discharging during the respective operations and
b) State of Charge (SOC), that defines the percentage of charge in the battery of the connected EV.

Hence, we need to inspect packets to monitor these quantities. During charging and discharging operations, the EV periodically fetches the flow reservation list from the grid while also periodically updating its power status to the grid. The flow reservation list contains the start and end of the subscription period. This data is used to verify that there are no power status updates outside the specified time interval. Power status updates occur only during the subscription period. This component of the anomaly detection engine handles attack 4 in Tab. I. The power status update messages contain vehicle SOC related information in terms of amount of power drawn. These power measurements are validated against physical power measurements as discussed next.

D. Consistency of Cyber and Physical States

We develop additional security checks to verify the consistency of cyber states with the physical states. Power measurements related to the EV charging load can be obtained from the sensors installed on the power distribution grid. Potential sources for the actual power data include the connected EV charging equipment, smart meters and distribution line measuring devices. These sensors are assumed to be able to report the instantaneous active power or even complex power readings that are time-stamped for validating the cyber states of EV activities.

To this end, let $P(t)$ denote the active power measurement for time period t, with the time resolution of Δt. We plot a sample daily profile of $P(t)$ for an actual residential home in Fig. 5, both with and without the EV charging load. The residential home power demand and the EV charging data are obtained from the Pecan Street database at a minute-level sampling rate [1]. As observed from the data, the EV charging load exhibits unique pattern that is different from other residential appliances and devices. First, the EV power demand is at least 3kW and also higher than that of other typical household loads [5]. Based on the observed power profiles from the Pecan Street dataset, only the air-conditioning load has a comparable level of power demand. Second, the EV charging typically lasts for hours at the constant rated power demand level, which is different from the periodic pattern of air-conditioning load. Therefore, for the EV load, there is a noticeable change only at the start and end time points of charging.

To better demonstrate this feature of EV load, we process the power data using a simple high pass filter to determine the rate of change between two consecutive data points, namely $\Delta P(t) = P(t) - P(t - \Delta t)$. To capture the fast change due to EV charging, the sampling rate of the power data needs to be sufficiently high to show that the EV can reach its rated power within two minutes while all other residential loads stay relatively unchanged. If the sampling rate increases, one may need to perform more sophisticated filtering process to recover the EV charging states. The filtered output for the residential home load with EV charging in Fig. 5 from approximately 3:35 – 6:00 AM is illustrated in Fig. 6. Note that the negative power in the aggregated load in Fig. 5 is due to PV generation. The first spike in the plot represents the time instant when the EV charging started, reaching its rated power of around 3kW within two minutes. Note that when the charging ended for this 3kW rated EV, its power consumption slowly drops from the rated power and is therefore not noticeable from the filtered output. However, for EVs rated around 6kW or higher from the Pecan Street dataset, the end charging time is more noticeable from the filtered output. Specifically, the 'end charging' characteristics are very similar to the 'start charging' ones as the EV's power consumption drops from its rated power to zero within two minutes. This will result in a negative spike in the filtered output with a magnitude close to its rated power.

The power data spike due to start/end of EV charging is used to verify the physical state when a packet with power measurement message is detected on the cyber network. Note that if the load profile is at sufficiently high sampling rate, it is unlikely that there is other major load change activity at the start/end time of EV charging. Therefore, the turn-on/-off events of air-conditioning loads will not confuse the engine with a potential EV activity. If the sampling rate gets slower, it will be necessary to smooth out the non-EV loads in order to determine the power spike from EV charging. Note that this may reduce the confidence in physical state verification part due to the existence of other heavy loading appliances with periodic patterns (*i.e.,* air-conditioning load). This component of the anomaly detection engine handles attacks 1,2 in Tab. I.

We have performed a design space exploration for many combinations of EV and household numbers for the SAE protocol family. For instance, *(i)* it can handle multiple EVs at a single household even if they start charging at the same instance of time. In which case, the total power consumed by these EVs is used to validate consistency on cyber and physical sides *(ii)* it can can detect even if an EV stops charging partway before reaching 100%. Due to space constraints, we do not present all the details here but in the online tech report [14].

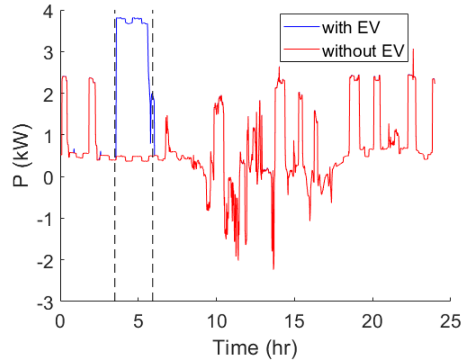

Fig. 5: Load profile of a residential home with and without the EV charging load.

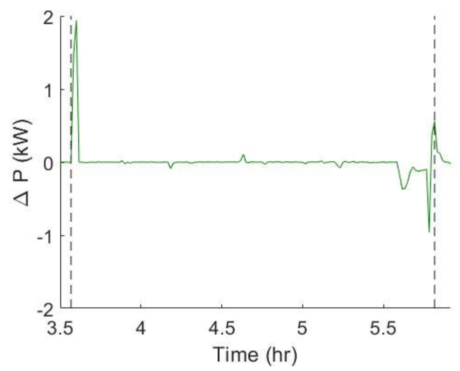

Fig. 6: Filtered output of the total EV included load in Fig. 5 during the EV charging period.

V. Evaluation

To evaluate the anomaly detection engine, we implemented a prototype in Python2.7 on the Intel i7 NUC platform with specifications as follows:

- Platform - Intel i7 NUC
- Processor - Intel(R) Core(TM) i7-7567U CPU @ 3.5 GHz, 4 cores
- Memory - 32 GB RAM, 128 GB HDD
- Operating System - Ubuntu 16.04

Our goal is to not only detect anomalies accurately but to also do so as soon as a malicious network packet arrives at the aggregator. The anomaly detection engine is placed at the aggregator in the V2G system as shown in Fig. 2. This makes it important to ensure that it does not introduce significant delay to the packet transfer rate at the aggregator. We therefore evaluate the prototype of our anomaly detection engine in terms of *(i)* **Accuracy** is measured in terms of *false positives* and *false negatives*. *(ii)* **Performance** is measured in terms of *average* and *worst case* time taken by the anomaly detection engine to inspect one packet and compared with minimum network latency. Note that as mentioned previously, since there are no existing anomaly detection engines to the best of our knowledge that combine cyber and physical properties along with communication standards for this type of system (§VII), direct comparisons are not possible.

Power Data for EV Charging - Packets on the cyber network are generated based on EV power profile data for charging obtained from the Pecan Street database at a minute-level sampling rate [1]. This data consists of timestamps and corresponding power measurements with respect to each EV, from which all required information (such as subscription period) for packet generation can be extracted.

Power Data for EV Discharging - Using the above EV charging data, we formulated the data for EV discharging based on efficiency formulas [6]. From [[6], eq. (12)], the magnitude of EV charging/discharging power from/to the power grid is described as follows

$$|P_{grid}| = \begin{cases} \dfrac{EV_{rated}}{\eta_1} = |P_c| \\ EV_{rated} * \eta_2 = |P_d| \end{cases} \quad (1)$$

EV_{rated} is the rated EV power, η_1 is charging efficiency, η_2 is discharging efficiency, $|P_c|$ is the magnitude of EV charging power and $|P_d|$ is the magnitude of EV discharging power. Using algebraic manipulation and substituting EV_{rated}, we obtain the magnitude of EV discharging power as a function of EV charging power. From the reference paper [6], we obtain the charging and discharging efficiencies in Table II as 0.92 and 0.92 respectively. Plugging these values into the previous equation, we get the magnitude of the EV discharging power is approximately 85% of the EV charging power.

$$|P_d| = |P_c| * \eta_1 * \eta_2 = |P_c| * 0.92 * 0.92 = |P_c| * 0.846 \quad (2)$$

Currently, there are no EVs/EVSEs that support the SAE J2847/1 and SAE J2847/3 standards since these communication standards are still in the process of development. There has been significant effort towards making the real world implementation of SAE standards feasible [16], [17]. Once these standards are implemented, our anomaly detection engine can be easily integrated with them. Therefore, to test our anomaly detection engine, we generate packets with custom HTTP payloads according to specifications provided by SAE communication standards [8]. Further details on how the packets are generated are provided below.

A. Accuracy

Testcases to test for false positives are generated based on certain ground truth and testcases to test for false negatives are obtained by modifying the former testcases to introduce various kinds of anomalies.

Message Sequence Validation - The ground truth for testcase generation here is the SAE standards [8]. *(a)* Testcases for false positives include all possible valid sequences consisting of parallel as well as repeating sequences. Due to the possibility of a lot of valid variations, there is a large number of these test cases. For instance, consider the charging operation. As shown in Fig. 4, first a flow reservation with or without cancellation is performed (note the existence of two possibilities already). Then the EV starts sending periodic power updates to the grid during its subscription period for charging. In parallel, the EV also periodically fetches the flow reservation list from the grid. Periodic messages give rise to repeating sequences and increase the number of possible valid variations. Similarly, parallel sequences of messages (power updates and fetching of flow reservation list in this case) also increase the number of possible valid variations. This is because one or more messages from a parallel sequence (say, power updates in this case) can arrive anywhere between messages in a related parallel sequence (fetching of flow reservation list in this case). The sequence to which the message belongs is identified using the message data. *(b)* Testcases for false negatives were generated by randomly placing invalid packets amidst valid sequences.

Message Frequency Validation - The ground truth for testcase generation here is again the SAE standards [8]. *(a)* Testcases for false positives consisted of packets with expected message periodicities. *(b)* Testcases for false negatives consisted of packets containing messages with periodicities different (*i.e.,* periodicities lower and higher than expected values) from expected values as specified in the SAE standards [8].

Subscription Period Validation - The ground truth here is based on the data obtained from the Pecan Street database [1]. *(a)* Testcases for false positives consisted of packets with subscription periods consistent with the aforementioned data. *(b)* Testcases for false negatives consisted of packets with inconsistencies where the actual subscription period was different from previously agreed upon subscription period, i.e, the arrival time of packets containing power status updates were modified so as to be different from the time intervals specified in packets containing the flow reservation list.

Consistency of Cyber and Physical States - We have a high certainty on detecting the stop charging time for EVs with higher rated power versus EVs with lower rated power for two reasons. The first is that the stop charging characteristic for higher rated EVs is very similar to the start charging characteristics of EVs except that it will drop from its rated power to zero within two minutes. The lower rated EVs take longer than two minutes to stop charging so there will not be a noticeable negative spike in the filtered load sequence. The second reason is that the higher rated EVs are less sensitive to non-EV loads changes. For instance, if an 1.5kW AC unit starts/stops around an EV start/stop event, it will effect the total ΔP of two consecutive time stamps. In the worst case scenario, the AC will reach its rated power within one minute. The AC rated power is 50 percent of an EV rated at 3kW which will force ΔP out of the acceptable range of start/stop charging values. The AC rated power is only 25 percent of an EV rated at 6kW, however, ΔP will usually be barely within range of acceptable start/stop charging values making higher rated EVs less sensitive to non-EV load changes. The ground truth here is again based on the power data obtained from the Pecan Street database [1]. *(a)* Testcases for false positives consisted of packets with power measurements consistent with this data. During steady state charging, the EV power should not vary beyond +/- 0.5kW in the worst case scenario. Any power measurement beyond this range is considered anomalous. *(b)* Testcases for false negatives makes use of this fact, *i.e.,* it consists of packets with inconsistencies with respect to power measurements, i.e, power measurements were varied to be above and below this permissible range of +/- 0.5kW.

B. Performance

Minimum network latency - According to the smart grid communication requirements specified by the Department of Energy, the minimum network latency with reference to Electric Transportation is *2 seconds* [4]

Worst case time taken - We compare the worst case time taken by the anomaly detection engine to inspect a packet with the minimum network latency to determine whether or not the anomaly detection engine introduces significant delay. The worst case time taken to inspect a packet is when it goes through the longest datapath in the anomaly detection engine. The 99.9th percentile worst case time is *0.165 seconds*.

Average time taken - In this evaluation, the total number of EVs that are simultaneously handled by the aggregator has been varied (up to a maximum of 400) based on literature [11]. As shown in Fig. 2, there are multiple households connected to a single aggregator with (possibly) multiple EVs connected to each household. The power measurements originally obtained from the Pecan Street database [1] contains individual EV loads as well as the corresponding aggregated loads at households where one household has just one EV associated with it. We have modified this power data as follows: *(i)* In order to simulate multiple EVs connected to each household, we first assign charging/discharging start times to include situations where multiple EVs at one household start to charge/discharge either at the same instance of time or at different instances of time. Then, we compute the operation end times while keeping the duration of operation unchanged. Note that we are only modifying the charging/discharging start times for EVs *i.e.,* we are just sliding the EV charging/discharging duration windows. Algorithm 7 on data generation can be found in the online tech report [14]. *(ii)* To simulate multiple such households being simultaneously connected to the aggregator, we replicated the power data per household. The average time taken by the anomaly detection engine to inspect a packet is approximately *0.014 seconds* as shown in Fig. 7. Tab. II summarizes these results. As stated previously (§II), we assume that the cyber network connecting physical network sensors to the anomaly detection engine has similar properties (in terms of bandwidth etc) as the cyber network in the V2G system. Therefore, the individual EV power measurements on the V2G cyber network and the corresponding aggregated power measurements on the sensor cyber network are received by the anomaly detection engine simultaneously (*i.e.,* without any delay).

TABLE II: Performance of Anomaly Detection Engine

Minimum Network Latency	**Worst case per packet inspection time**	**Average per packet inspection time**
2 seconds	0.165 seconds	0.014 seconds

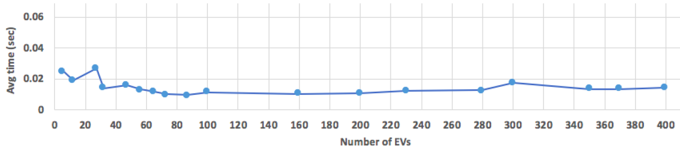

Fig. 7: Performance of the Anomaly Detection Engine

VI. Discussion

As seen from the results in Tab. II, the worst case time to inspect a packet *i.e., 0.165 seconds* is lesser than the minimum network latency, *i.e.,* 2 seconds. Therefore, our anomaly detection engine detects malicious packets almost instantly without introducing significant delay at the aggregator.

In all of our validation techniques, the consistency of data obtained from one source is verified against data obtained from a different source as explained next. This makes it more difficult to tamper with the data so as to produce false consistencies. In case of Message Sequence Validation and Message Frequency Validation, packets are verified against information obtained from the communication standards documentation. For validating consistency of Cyber and Physical States, power data reported in cyber messages is validated against physical power measurements. In case of Subscription Period Validation, subscription period related data is validated using messages from the power grid side of communication that contain flow reservation list and messages from the EV side of communication that contain power status updates.

Packet loss can result in certain packets either arriving (at the aggregator) out of sequence or at unexpected times. This might result in false positives being signalled by the anomaly detection engine. However, rate of packet loss is very low given the 99-99.99% network reliability requirement for Electric Transportation in the Smart Grid [4]. Network delays/jitters could also lead to false positives due to packets arriving at an unexpected time. This can be handled by setting an appropriate tolerance level with respect to time constraints. Our anomaly detection engine has a complete list of all possible valid states and state transitions. This eliminates the likelihood of false negatives. Our evaluation is carried out based on the previously stated assumption that the inputs obtained from atleast one of the cyber or physical networks are genuine. If this is not true, then an attacker could tamper with the cyber and physical states of the system to obtain consistency and conceal the attack from our anomaly detection engine.

VII. Related Work

Securing the advanced metering system by using specification based intrusion detection [2] monitors just the cyber state of the system by observing traffic among access points and meters at various layers to ensure expected behavior. We use similar techniques to monitor the cyber state of a V2G aggregator. In addition, we also check for consistency of cyber and physical states of the system. The paper by Liao et al [12] focuses on enhancing power grid security by using micro-synchrophasors as a tool to monitor and manage distribution networks. This work is similar to our work in that it uses data from sensors for monitoring. However, in addition to using sensor readings, our anomaly detection engine also monitors other data constraints related to the communication protocol specifications along with timing constraints related to message frequencies and subscription periods. Chen et al [3] propose an efficient and secure authentication scheme for V2G networks that preserves privacy. The paper focuses on securing the communication of EVs in the V2G system of power grid. On the other hand, our work focuses on securing the aggregator, an important component of the V2G system, by increasing its resiliency to attacks. There is some related work on identifying EV charging profiles for improving power distribution system operations. The statistical characteristics of EV's state-of-charge or the duration of charging period have been studied in [18], [20] by analyzing a fleet of EV charging profiles. More recently, a deep learning approach has been proposed in [21] to extract the EV charging profile from the aggregated household demand as a load disaggregation problem. However, we use a similar filtering mechanism but for a different purpose, *i.e.,* anomaly detection to ensure security. We have developed a simple approach for estimating EV charging status that is very suitable for real-time implementation needs of the proposed anomaly detection engine.

VIII. Conclusion

In this work, we have presented a novel architecture of a Cyber-Physical Anomaly Detection Engine that captures the cyber and physical properties of the system along with the related communication standards to define correct system behavior. The simple model of our anomaly detection engine demonstrates that accurate and almost instantaneous detection of anomalies is feasible. Although our prototype Cyber-Physical Anomaly Detection Engine is based on SAE standards of communication for V2G system, this architecture can be extended to other communication standards for other DERs as well.

Acknowledgements

This work is supported in part by grants from the Dept. of Energy (CEDS program - contract number DE-0E0000826) and the National Science Foundation (NSF SaTC 1718952 and ECCS 1902319). Any opinions, findings, and conclusions or recommendations expressed here are those of the authors and do not necessarily reflect the views of sponsors. The authors would also like to thank Tim Yardley, Steve Granda, Devu Shila, Lynn Ren, Tom Carroll, Penny Mckenzie and Jereme Haack for their useful discussions and ideas.

References

[1] https://dataport.pecanstreet.org/.

[2] R. Berthier, W. H. Sanders, and H. Khurana. Specification-based intrusion detection for advanced metering infrastructures. In *PRDC*, 2011.

[3] J. Chen, Y. Zhang, and W. Su. An anonymous authentication scheme for plug-in electric vehicles joining to charging/discharging station in vehicle-to-grid (v2g) networks. In *China Communications*, 2015.

[4] DOE. Communications requirements of smart grid technologies. 2010.

[5] H. Estes, S. Santoso, and G. Fisher. Analysis of high-resolution electric vehicle charging on time-of-use grid demands. In *IEEE Power & Energy Society General Meeting*, 2015.

[6] L. Gkatzikis, I. Koutsopoulos, and T. Salonidis. Decentralized vehicle-to-grid control for primary frequency regulation considering charging demands. In *IEEE Transactions on Power Systems*, 2013.

[7] L. Gkatzikis, I. Koutsopoulos, and T. Salonidis. The role of aggregators in smart grid demand. In *IEEE Journal on Selected Areas in Communications*, 2013.

[8] SAE International. What do sae standards offer the automotive electronics industry? 2009.

[9] W. Kempton and J. Tomic. Vehicle-to-grid power fundamentals: calculating capacity and net revenue. In *Journal of Power Sources*, 2005.

[10] W. Kempton and J. Tomic. Vehicle-to-grid power implementation: from stabilizing the grid to supporting large-scale renewable energy. In *Journal of Power Sources*, 2005.

[11] A. Y.S. Lam, K.C Leung, and V. O.K. Li. Capacity estimation for vehicle-to-grid frequency regulation services with smart charging mechanism. In *arXiv:1410.1282v4 [cs.SY]*, 2015.

[12] A. L. Liao, E. M. Stewart, and E. C. Kara. Micro-synchrophasor data for diagnosis of transmission and distribution level events. In *Transmission and Distribution Conference and Exposition, IEEE/PES*, 2016.

[13] Nescor. Electric sector failure scenarios and impact analyses version 3.0. 2015.

[14] C. Niddodi, S. Lin, S. Mohan, and H.Zhu. Secure integration of electric vehicles with the power grid. In *arXiv:1905.01035v4 [cs.SY]*, 2019.

[15] C. Niddodi, S. Mohan, and T. Yardley. Securing electric vehicles in the power grid. In *Proceedings of 29th Euromicro Conference on Real-Time Systems Work-in-Progress Session*, 2017.

[16] R.M. Pratt, F.K. Tuffner, and K. Gowri. Electric vehicle communication standards testing and validation phase 1: Sae j2847/1. 2011.

[17] R.M. Pratt, F.K. Tuffner, and K. Gowri. Roadmap for testing and validation of electric vehicle communication standards. In *Electric Vehicle Symposium (EVS26)*, 2012.

[18] K. Qian, C. Zhou, M. Allan, and Y. Yuan. Modeling of load demand due to ev battery charging in distribution systems. In *IEEE Transactions on Power Systems*, 2011.

[19] T. Seals. Electric vehicle charging stations open to iot attacks. 2018.

[20] W. Su and M.Y. Chow. Performance evaluation of an eda-based large-scale plug-in hybrid electric vehicle charging algorithm. In *IEEE Trans. Smart Grid*, 2012.

[21] S. Wang, L. Du, J. Ye, and D. Zhao. Robust identification of ev charging profiles. In *IEEE Transportation Electrification Conference and Expo*, 2018.