Detecting Energy Theft Attacks on an Off-Grid Charging Station

Anatolii Khalin
CentraleSupélec
CESSON-SEVIGNE
ILLE-ET-VILAINE, France
anatolii.khalin@centralesupelec.fr

Jean-François Lalande
CentraleSupélec
CESSON-SEVIGNE
ILLE-ET-VILAINE, France
jean-francois.lalande@centralesupelec.fr

Romain Bourdais
CentraleSupélec
CESSON-SEVIGNE
ILLE-ET-VILAINE, France
romain.bourdais@centralesupelec.fr

Abstract

With the rapid development of charging infrastructure for Electric Vehicles, the risks of cyber-physical attacks, including energy theft are growing. The attack detection results of energy theft are usually validated on real open access data of charging sessions, however, the attacks themselves are artificially introduced. To address this issue, this work presents a three-week experiment on a real testbed including the production and consumption of energy with realistic energy theft attacks occurring in the system. The energy setup emulates a charging bike station where users can charge bikes at different levels of state of charge and at different durations of charging sessions. The attacker is one of the users who steals energy from the system for its own bike and can override the reported consumption of power. We propose a method for detecting such attacks based on the total production/consumption power balance. The full dataset of the three-week experiment is published with this work for reproducibility purposes.

ACM Reference Format:
Anatolii Khalin, Jean-François Lalande, and Romain Bourdais. 2025. Detecting Energy Theft Attacks on an Off-Grid Charging Station. In *The 16th ACM International Conference on Future and Sustainable Energy Systems (E-ENERGY '25), June 17–20, 2025, Rotterdam, Netherlands.* ACM, New York, NY, USA, 7 pages. https://doi.org/10.1145/3679240.3734650

1 Introduction

Research on the cybersecurity of Electric Vehicles (EV) charging has been especially extensive over the last decade. Related surveys deal with the security of charging protocols and the threats the technology faces [8, 11–13]. The considered threats include the problem of energy theft or unauthorized charging. It can be considered as the threat that will probably grow and become a challenge with the widespread adoption of Electric Vehicles [3, 9].

Most of areas are currently covered by public, residential, or semi-public charging stations, powered by a public electric grid (power plants). However, the trend of installation of off-grid charging stations is going upward, due to some clear advantages: accessibility in remote locations and clean energy power sources. In addition, due to the high energy demand of new electric vehicles [4, 25], burdens on the public grid can occur, which suggest that stations should include off-grid capacities. For instance, the so-called FEVER initiative [17] was introduced to decouple EV charging demand from the electricity grid in the UK, also facilitating a more rapid deployment of charging infrastructure. Additionally, for areas with poor grid connectivity (for example CHARGE Solar-Powered Stations in South Africa[1]), off-grid capabilities become mandatory. An extensive review on off-grid and hybrid charging stations can be found in [24], concluding in an overwhelming popularity of hybrid systems over off-grid ones, with predominant energy sources combining PV and batteries.

There are key differences between energy theft from charging stations and the public grid. The most common practice of stealing energy from the public grid is a physical bypass, such as illegal power line tapping [18]. In such a case, only one user is supposed to consume energy and would be easily spotted as an attacker in case of inconsistencies in the measured consumption. On charging station, various users are involved and the technologies involved can have flaws for example in the radio protocols [6] or authentication mechanisms [2]. Additionally, the hardware can be physically attacked by modifying the involved sensors or measurement systems. If an attack succeeds on a charging station, it can be replicated over all stations involving the same setup. For these reasons, we focus on the specific case of an attacker stealing energy from a charging station, especially since such attacks can be extended to a public grid and cause major issues for the infrastructure [1].

Despite the plethora of openly available data sets on electric transport charging data (Boulder city[2], Palo Alto city[3], Netherlands[4], Dundee city, Perth & Kinross Council, etc.), there is no dataset including real attacks during normal operations of a charging station. Researchers are forced to inject anomalies into existing experimental data artificially [5, 6, 16] which can mislead the detection results if the injected events contain too obvious differences with normal events. Using artificial intelligence methods, the detection of anomalies can give over estimated results, which justify to perform attacks during the real running of the system to guarantee that the data that is similarly recorded for legitimate and malicious events.

In this paper, we design a testbed representing a bike charging station composed of real solar panels, batteries and grid connectivity. We run the testbed and consume energy corresponding to predetermined arrival of users with their bikes. Bikes are simulated with charging profiles and operated with a programmable device

This work is funded by Direction Générale de l'Armement through CREACH LABS under the project CAMTAR.

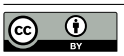

E-ENERGY '25, Rotterdam, Netherlands

ACM ISBN 979-8-4007-1125-1/25/06
https://doi.org/10.1145/3679240.3734650

[1] https://charge.co.za/
[2] https://open-data.bouldercolorado.gov/datasets
[3] https://data.cityofpaloalto.org/dataviews/257812/electric-vehicle-charging-station-usage-july-2011-dec-2020/
[4] https://platform.elaad.io/download-data/

consuming the energy from the system. During a one week running of the system, an attacker charges sporadically his bike by consuming energy and removing the reported consumption from the system. The attacker model is similar to the one introduced in [7]: the attacker can corrupt by physical means or remotely the control system and modify one of the sensor values. We run two experiments: during one week the attacker charge his bike with high levels of current, and in a second one, the levels are lower. For both experiments, the solar panels, the batteries are involved, and eventually the grid if the state of charge of batteries becomes insufficient. We introduce a detection algorithm based on the correlation of the values of the system in order to estimate the intervals of time where an attack occurs. Our results show that about 78.92% (week 2) and 65.53% (week 3) of stolen energy is detected in average with almost no false positives reported. Additionally, we released our dataset[5] containing all measured values of our testbed for further usage.

The outline of the paper is as follows. Section 2 provides related works about energy theft detection and security of electric charging. Section 3 is describing the scenario that is simulated on our testbed, and the model of attack. Section 4 gives the detection performance of the proposed detectio algorithm. Finally, Section 5 concludes the paper and gives future directions of the project.

2 State of the art

The problem of energy theft has been addressed in the literature mostly in the context of public power grids. In [21], the issue of stealthy energy fraud in power grids due to the increase in the use of corruptible smart meters has been investigated. The authors provide two data mining techniques to identify energy theft. A smart meter testbed in the University of Illinois at Urbana-Champaign was used as an experimental validation for representing the residential area power grid. Realistic load profiles based on simulated residents were generated for the experiment. In [15], the authors proposed a method of a consumption pattern-based energy theft detector, which compares the total consumption measured by transformer meters to the total usage reported by smart meters. Authors trained an SVM on a public anonymized dataset of 5000 real customers. Later, in [14], the same problem was addressed with a CNN-LSTM-based deep learning technique. The method was evaluated with the real data of 9655 users, provided by State Grid Corporation of China. In [20], the coordinated pricing and energy theft attack in the Smart Home Cyber-Physical System was considered. The method of partially observable Markov decision process was used to select non-compromised households and to fix or disregard information from the corrupted ones. The validation consisted of a simulation, including the Smart Home model. In [23] feature engineering and machine learning algorithm was proposed, with a validation on data of more than 4000 households data. With a similar goal to these works, we present an energy theft attack on a real energy-based testbed, where the target of the attack is a smaller-scale local grid rather than a public one.

Despite the recent advances of the EV and other electric transport infrastructure, there appears to be only a few results directly addressing the issue of energy theft in the literature. For instance, in [5] the price false data injection on a hybrid charging station using PV was considered. A combination of a statistical Greedy Gaussian Segmentation (GGS) for anomaly detection with kNN for its source identification is introduced. For the algorithm validation experimental data of real EV charging with 7994 records from the UCLA campus was used and the attack was artificially injected. In [6] the experimental datasets from multiple open sources were artificially extended with anomalies related to attacks on energy consumption values and with errors, which leads to energy theft among other consequences. A collaborative Anomaly Detection System was introduced based on local retraining of ML models based on collectively available data. The results presented in this paper aim to detect energy theft attacks on a charging station based on measured consumption readings like in [6], but using a static algorithm rather than complex machine learning methods.

In [7], the so-called overcharging covert attack on an Electric Vehicle Supply Equipment is introduced: the physically tempered increase in power consumption by the battery of an EV is canceled by an FDI (False Data Injection) attack measurement. The detection strategy includes a dynamical model of the battery, based on which two detection algorithms are provided: a static one, which compares the measured battery voltage with the predetermined upper limit; and the dynamical which compares a measured battery voltage with an expected one, based on the model from the Fault Detection and Identification techniques. The results are then compared and validated on a charging simulation with artificially injected noise. The results are later extended in [22] with new adversarial threat models. In our work, similarly to [7], we simulate attacks using False Data Injection, but we inject attacks in real-time to get more realistic detection results and to better represent the impact.

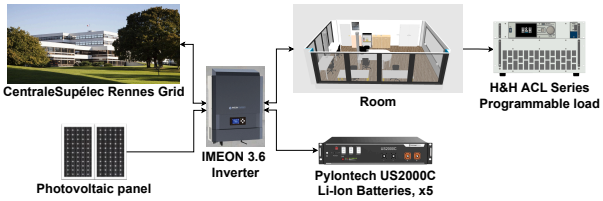

Figure 1: Testbed

3 Testbed and attack scenario

This section presents a well-monitored energy-based testbed for hosting scenarios related to cybersecurity and energy management. The testbed is used to deploy a scenario where an attacker is stealing energy from a charging station for bikes. We first describe the testbed and then discuss the attacker's model and how we simulate the charging station and the attacks in the testbed.

3.1 Testbed description

As shown in Figure 1, the testbed is hosted in a room where an inverter is the central link with the grid of the campus, the photovoltaic panels, Li-Ion rechargeable batteries and a programmable load. The inverter (IMEON 3.6) distributes the energy as follows: all the energy generated by the PV panel goes directly through the

[5]https://doi.org/10.5281/zenodo.15297511

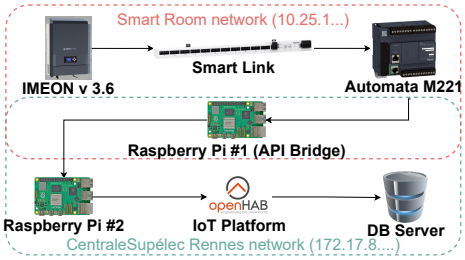

Figure 2: Data collection

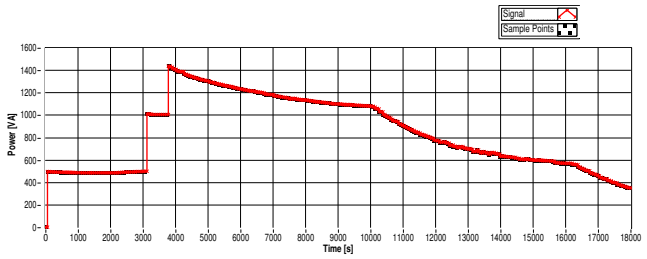

Figure 3: Daily charging profile (3 bikes)

Smart Inverter, which decides the destination for the storage or further transmission. Based on the mode of operation, the energy from PV can be either stored in the batteries, or transferred directly to the campus grid. By design, the Inverter IMEON allows switching between four modes of operation: Smart Grid (priority on PV and batteries power utilization over grid), Back-Up (priority on the full charge of the batteries from PV and grid), On-Grid (the battery is not utilized), and Off-Grid (island mode with no public grid utilized). For the performed experiment described later in Section 4, we use a Smart Grid mode to simulate a hybrid system *i.e.* that order the consumption first from the PV, then from the batteries and finally from the campus grid. However, during the night when the renewable energy is not available, we switch to Back-Up mode to charge the batteries from the grid, if necessary. The batteries that store the energy collected by the PV panels are 5 Li-Ion Pylontech US2000C batteries, located in the thermal chamber and connected in parallel with a total nominal capacity of 12 kWh. The H&H ACL Series Programmable load is connected to the local AC grid and can consume energy from all sources, through the IMEON inverter.

Figure 2 shows the network topology of the testbed. The Smart Room network (10.25.1.x) connects the IMEON inverter to an automata and can be interrogated from a secondary local network (172.17.8.x) through a raspberry Pi #1. The raspberry Pi #1 plays the role of a bridge between the two networks and offers an HTTP API to read the register values of the IMEON using the modbus protocol. Such a setup is representative of industrial control systems where low level networks transport modbus messages, without any security involved, and where high level networks contain web applications for supervision purpose.

The Raspberry Pi #2 hosts the open-source *OpenHAB* software for tracking in real time all energy values and storing them in a database. *OpenHAB* also allows to write orders into the inverter, such as the changing of operation mode. In particular, the values collected for our scenario can be expressed in four groups:

$$\text{I} : \{P_{\text{PV}}, I_{\text{PV}}, V_{\text{PV}}\}, \qquad \text{II} : \{I^{\text{ch}}_{\text{batt}}, I^{\text{disc}}_{\text{batt}}, V_{\text{batt}}, SoC\},$$
$$\text{III} : \{P_{\text{grid}}, I_{\text{grid}}, V_{\text{grid}}\}, \quad \text{IV} : P_{\text{load}}.$$

- the first group represents simultaneous measurements of the power production, current, and voltage of the PV panel;
- the second group consists of the measurements from batteries (charging current, discharging current, voltage, and state of charge);
- the third group provides values from the local AC grid: the power readings go to negative when the system produces an overhead of energy and sends it to the public grid, else it is positive;
- the final group is dedicated to the power consumption readings of the electric load, which in our case is assumed to be attacked as we will explain in the next section.

3.2 Scenario: bike charging station

We chose to deploy real attacks in a scenario where users charge their bike at a charging station that is powered by our three sources. The scenario has been chosen to obtain the more realistic data given the setup in place.

- The station can charge up to 3 bikes at a time;
- The primary source of charging is the energy coming from the photovoltaic panels;
- The secondary source of charging is the Li-Ion batteries;
- The batteries are recharged from the public grid at night if necessary.

In order for the experiment to be as realistic as possible when a bike is charging we used charging profiles from the project *WeBike* of University of Waterloo [10]. During this project, the 33 electric bikes equipped with various sensors were distributed to the faculty and all the readings were recorded in a dataset[6]. In particular, we extracted the bike charging profiles of June and May 2016 and adapted them to command the programmable load of our testbed. In particular, we performed the following operations:

- Fixing the time steps to 60 seconds;
- Eliminating zeros or other obvious errors;
- Scaling the data to fit the Smart Room power consumption limit (3000 W).

In order to represent the load of a daily venue of three bikes, we combined the profiles by scheduling the arrival of bikes 1 to 3, resulting in the total consumption profile of Figure 3. The power consumption profile of Figure 3 represents a three bike charging session, with only one bike charging at first, following by a second after a delay of about 50 min, and the third one following with a delay about 65 min. Each bike is represented by a standard CC-CV (Constant Current/Constant Voltage) protocol, which leads to asynchronous drops in maximum power (as can be observed around 166 and 270 minutes).

[6] https://borealisdata.ca/dataset.xhtml?persistentId=doi:10.5683/SP2/7OAETS

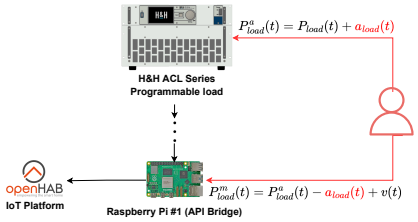

Figure 4: Attack Implementation

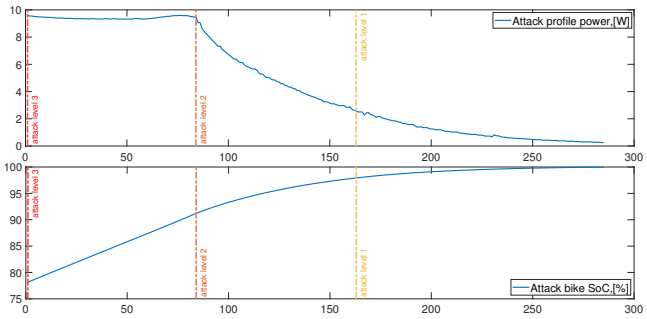

Figure 5: Attack levels

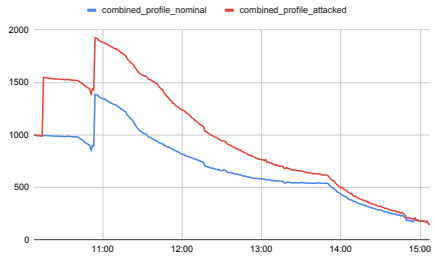

Figure 6: Attack profile, levels 1-3

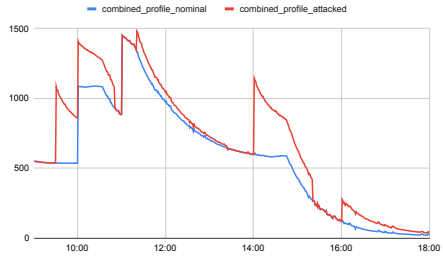

Figure 7: Attack profile, levels 1,2

The regular operation of the charging station (*i.e.* without attacks) includes 5 days of charging sessions, where energy is provided by PV panels and batteries. Each day, new profiles of either 2 or 3 bikes were charging with scheduled arrival delays.

3.3 Attack model

The energy theft attack was designed as a combination of a physical and integrity (FDI) attack. The idea is that the adversary pulls the electricity for the bike when he is not allowed to (such cases are prominent in practice, mostly in China [26], and sometimes lead to fire accidents). Furthermore, in order to avoid being discovered, our attack model assumes that the attacker has the ability to gain access to the API bridge (due to the software vulnerability) and to override the energy consumption that is reported from the inverter. As a consequence, such attacker has strong power over the system because we suppose that he is able to perform FDI attacks, similarly to [7], and subtract its power consumption in real time, according to the power of his bike. As a summary, we perform in the testbed the following operations when the attacker bike is charging:

- (Physical): Increase the requested power by sending a corrupted profile that adds a_{load} to the programmable load;
- (FDI): Subtract the increased power a_{load} every data reading time step in the API bridge (which emulates a cyber attack similar in the principle to the StuxNet attack [19]).

As represented in Figure 4, the control signal sent to the programmable load, denoted by P^a_{load} is then applied as a consumption and measured by the IMEON, including the measurement noise $v(t)$. On the other side, a new event P^m_{load} is sent to the API bridge with a subtracted FDI attack a_{load}.

In order to introduce diversity in attack's behaviors in our scenario, we defined three different levels of power for the attacker. Three stages of Constant Current/Constant Voltage (CC-CV) charging are represented in Figure 5: the top part of the figure shows a current profile of a bike charging session. The bottom part shows the state of charge of the bike's battery . The most noticeable part of charging (defined as attack of level 3) corresponds to the maximum power consumption, which corresponds to a charging up to ~90 percent of the charging capabilities. The next segment, from the peak of the current profile to 1/3 of the maximum defines the attack of level 2: it corresponds to charging from ~90 to ~96 percent of the battery. Finally, the last segment with the least power consumption is attributed to the attack of level 1 (from ~96 to 100 %). By mixing these levels of attacks and varying their duration, we introduce diversity in the difficulty of detection of these attacks.

Finally, we ran two weeks of experiments containing attack sessions:

- week 2: attacks that are triggered are full uninterrupted charging from level 3 to 1;
- week 3: attacks that are triggered contain only attacks of level 2 or 1.

Figure 6 shows the difference between a regular energy consumption (without any attack in blue: combined_profile_nominal) and the same consumption with an additional attack during a second week: the attacker consumes more energy during 5 hours, an additional 500 W at the beginning and almost 0 W at the end. For week 3, we show in Figure 7 the same comparison. Depending on the chosen level for the attack (1 or 2) and the time instance in this level, the difference can be more or less noticeable. For example, at 11:30, an attack of level 1 occurs until 14:00 with a low additional consumption. At 14:00 a new attack of level 2 is triggered but

Algorithm 1 Attack detection

Input: $P_e(1), \ldots, P_e(n)$
  for $i = 1, \ldots, n$ **do**
    if $|P_e(i)| > P_e^{\max}$ **then**
      $\delta \leftarrow \delta + 1$
    else
      $\delta \leftarrow 0$
    end if
    if $\delta > \bar{\delta}$ **then**
      for $j = 0, \ldots, \bar{\delta}$ **do**
        $\text{attack}(i - j) \leftarrow \text{true}$
      end for
    else
      $\text{attack}(i) \leftarrow \text{false}$
    end if
  end for
Output: $\text{attack}(1), \ldots, \text{attack}(n)$

stopped at 15:20. This example shows that an accurate detection can be more difficult in some cases when the attacker is stealing less energy.

4 Detection

In this section, we intend to evaluate our ability to detect attacks of weeks 2 and 3. We designed a detection algorithm based the measurement of error between the energy that is consumed and the energy that is produced. If the error is higher than a threshold, we suspect that an attack is occurring. Such algorithm is easy to setup and can run smoothly in the system. After presenting the algorithm, we evaluate the true/false positive/negative on the two weeks of experiments presented in the previous section.

4.1 Detection principle

We designed a detection algorithm that compares the combined consumption reported in openHAB and the combined energy production that we measure. The combined production P_{supply} can be expressed as follows:

$$P_{\text{supply}} = P_{\text{PV}} + I_{\text{batt}}^{\text{disc}} V_{\text{batt}} + P_{\text{grid}}. \quad (1)$$

Whilst, the total consumption by the equipments of the system P_{cons} can be expressed as follows:

$$P_{\text{cons}} = P_{\text{load}}^{m} + I_{\text{batt}}^{\text{ch}} V_{\text{batt}}. \quad (2)$$

The error P_e is introduced to accommodate for all the unmeasured consumption (such as consumption of energy by IMEON), disturbances (such as DC to AC conversion, PV measurement error), and measurement noises as follows:

$$P_e = P_{\text{supply}} - P_{\text{cons}}. \quad (3)$$

Based on the measurement of P_e, we can infer that an attack is occurring if P_e is too high. Thus, we designed a detection Algorithm 1 that iterates over the time and identifies if an attack is occurring at each time step i. In Algorithm 1

- $P_e(i) = P_{\text{PV}}(i) + I_{\text{batt}}^{\text{disc}}(i) V_{\text{batt}}(i) + P_{\text{grid}}(i) - P_{\text{load}}^{m}(i) - I_{\text{batt}}^{\text{ch}}(i) V_{\text{batt}}(i)$;
- $P_e^{\max}$ is the **error threshold** to be tuned according to the observed error of the first week of experiments;
- δ is the **number of consecutive readings considered as an attack**;
- $\bar{\delta}$ is a threshold that avoids false positives due to peaks of errors not linked to an attack;

Table 1: Energy consumption Experiment (Week 1)

Week 1 (kWh)	Total	PV + error	Battery	Grid	SoC (%)
Day 1	5.08	3.65	1.43	0	96
Day 2	5.13	0.63	4.5	0	37
Day 3	5.77	0.86	4.91	0	59
Day 4	4.18	0.74	3.44	0	71
Day 5	8.04	5.13	2.91	0	82
AVG	5.64	2.2	3.4	0	69

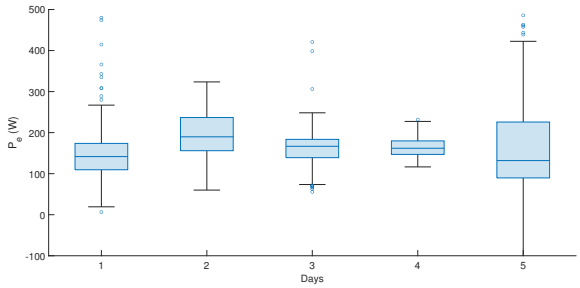

Figure 8: Error distribution, week 1

- $\text{attack}(i)$ is the final output of detection for each unit i of time.

4.2 Week 1 results: nominal operations

During week 1, the regular profiles (such as the ones of Figure 3) were sent to the programmable load to consume a power from the testbed, without any attack. Table 1 gives an overview of the energy consumption of the different parts of the testbed and the remaining state of charge (SoC) of batteries after each session. It can be seen from Table 1 that most of the consumed energy was provided by the batteries, due to the mostly cloudy weather (October-November), while PV was able to cover only about a third of the charging demand of 2-3 electric bikes. Nevertheless, for example in day 1, because of the sun, the SoC is high (96%) at the end of day. Figure 8 shows the distribution of P_e. We observe a persistent level of extra production of 160 W on average, which can be mostly attributed to the power consumption of the IMEON, which is not recorded by the supervision system and is considered as lost energy. Additionally, a significant number (111 out of 1721) of faulty error values are visible as outliers of the boxplot, presented in Figure 8. For example, during the days where the sun was available (day 1 and 5) the error outliers jump as high as 1200 and even produce a considerable amount of negative values. These errors can be explained by a non-optimal measurement of the PV power production, where the simultaneous power is measured once every 60 seconds instead of an average over the last minute. Such error can significantly impact the attack detection algorithm and can produce false positives.

Based on the observed error distribution of week 1, we fixed the two parameters of our detection algorithm: $P_e^{\max} = 320$ [W] and $\bar{\delta} = 3$ in order to limit false positives in the next experiments.

4.3 Week 2 results: level 3 attacks

Week 2 includes attacks of levels 3 to 1. We report in Table 2 the total energy consumption of the experiment for each day *i.e.* the energy consumed by users and the attacker. For example, on day 1,

Table 2: Energy consumption and attack detection (Week 2)

	Consumption (kWh)			Detection (%)					Attack			
Week 2	Total	Users	Attack	TPR	FPR	Accuracy	F1 score	Precision	Det.(kWh)	Det.(%)	Und.(kWh)	Und.(%)
Day 1	5.93	4.56	1.37	57.5	0	60.33	73.02	100	1.21	87.99	0.16	12.02
Day 2	5.12	3.75	1.37	47.33	0	53.17	64.25	100	1.06	77.74	0.31	22.26
Day 3	4.57	3.2	1.37	55.16	0	58.28	71.1	100	1.16	84.5	0.21	15.5
Day 4	4.75	3.38	1.37	43.42	0	47	60.55	100	1	72.8	0.37	27.2
Day 5	5.8	4.43	1.31	42.44	0	48.34	59.59	100	0.94	71.57	0.37	28.43
AVG	5.23	3.86	1.36	49.17	0	53.42	65.7	100	1.07	78.92	0.28	21.08

Table 3: Energy consumption and attack detection (Week 3)

	Consumption (kWh)			Detection (%)					Attack			
Week 3	Total	Users	Attack	TPR	FPR	Accuracy	F1 score	Precision	Det.(kWh)	Det.(%)	Und.(kWh)	Und.(%)
Day 1	7.51	6.32	1.19	40	0	56.32	57.14	100	0.86	72.46	0.33	27.54
Day 2	6.28	5.19	1.08	30.68	0	46.65	46.95	100	0.67	61.63	0.42	38.37
Day 3	5.82	4.78	1.04	41.14	0	62.38	58.3	100	0.74	70.87	0.3	29.13
Day 4	5.11	4.42	0.69	16.49	0	32.18	28.31	100	0.31	44.12	0.39	55.88
Day 5	4.94	3.76	1.18	43.13	1.55	57.61	60.04	98.74	0.93	78.56	0.25	21.44
AVG	5.93	4.9	1.04	34.29	0.31	51.03	50.15	99.748	0.7	65.53	0.34	34.47

the attacker consumed 5.93 − 4.56 = 1.37 Kwh. Then, we reported the detection performances: when we detect an attack at time i (TP: True Positive), when we do not detect an attack (FN: False Negative). To evaluate the performances, we computes the rates: TPR (True Positive Rate) = 100%·TP/(TP+FN), FPR (False Positive Rate) = 100%·FP/(FP+TN), Accuracy = 100%·(TP+TN)/(TP+TN+FP+FN), Precision = 100%·TP/(TP+FP), and F1 = 2·Precision/(Precision+TPR). We observe that, in average, 49.17% of all attacks have been detected and that no false positive is reported. This is due to the choice of the threshold $P_e^{\max}$ (see Figure 9). The associated quantity of energy that is reported as stolen is 1.07 Kwh *i.e.* 78.92% of the total consumed energy.

4.4 Week 3 results: level 1-2 attacks

The third week of experiments is reported in Table 3, where the overall detection drops to 34.29% (65.53% of stolen energy detected). This results show that the choice of low power attacks have a direct impact on the detection rate. Again, this detection ratio could be considered as low but it guarantees almost negligible FPR of 0.31%.

4.5 ROC curves

Figure 9 shows the ROC curves when the $P_e^{\max}$ varies, and the Figure 10 compares the same curves with different $\bar{\delta}$. ROC curves show that a higher detection rate implies a high ratio of false positive. For example, for week 3, detecting more than 70% of attacks would lead to 70% of false positives. This highlight the difficulty of the problem: at some point, the attacker is non distinguishable from errors or small variations due to the measurement components. Figure 10 show that when targeting low FPR (< 2%), values of $\bar{\delta} > 0$ gives better results. It can be explained by the fact that picks of error are not misleading the detection algorithm.

5 Conclusion and Future works

This paper presented a testbed for simulating a charging bike station where an attacker steals energy from the system. The strong hypothesis about the attacker capabilities is that he can control the reported power value in the system to cancel the power he steals from the system. We show how an algorithm can supervise and

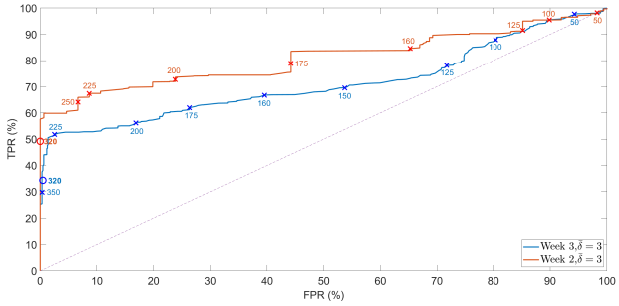

Figure 9: RoC for $Pe_{max} = 0, \ldots, 1000$, $\bar{\delta} = 3$, weeks 2,3

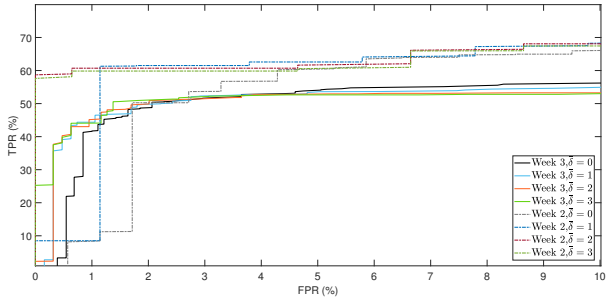

Figure 10: RoC for $\bar{\delta} = 0, \ldots, 3$, and $Pe_{max} = 0, \ldots, 1000$, weeks 2,3 (zoomed in up until $FPR = 10$.)

detect such attacks with decent detection results preventing energy theft attacks. The algorithm obtains acceptable results even when intermittent energy sources are involved, such as solar panels.

One of the future directions of the work is the improvement of the detection algorithm, with more extensive analysis of collected data. We also plan to extend this work with more complex scenarios, for example random decisions of the mode of the inverter, for instance, if an operator decides to stop the usage of batteries for other purposes. Another perspective of this work is to simulate

new types of attackers with other objectives such as preventing regular user to charge their vehicle.

References

[1] Samrat Acharya, Yury Dvorkin, Hrvoje Pandžić, and Ramesh Karri. 2020. Cybersecurity of smart electric vehicle charging: A power grid perspective. *IEEE access* 8 (2020), 214434–214453.

[2] Cristina Alcaraz, Jesus Cumplido, and Alicia Trivino. 2023. OCPP in the spotlight: threats and countermeasures for electric vehicle charging infrastructures 4.0. *International Journal of Information Security* 22, 5 (2023), 1395–1421.

[3] Cristina Alcaraz, Javier Lopez, and Stephen Wolthusen. 2017. OCPP protocol: Security threats and challenges. *IEEE Transactions on Smart Grid* 8, 5 (2017), 2452–2459.

[4] Alex Caines, Aritra Ghosh, Ankur Bhattacharjee, and Adam Feldman. 2021. The grid independence of an electric vehicle charging station with solar and storage. *Electronics* 10, 23 (2021), 2940.

[5] Yu-Wei Chung, Mervin Mathew, Cole Rodgers, Bin Wang, Behnam Khaki, Chicheng Chu, and Rajit Gadh. 2020. The framework of invariant electric vehicle charging network for anomaly detection. In *2020 IEEE Transportation Electrification Conference & Expo (ITEC)*. IEEE, 631–636.

[6] Jesus Cumplido, Cristina Alcaraz, and Javier Lopez. 2022. Collaborative anomaly detection system for charging stations. In *European Symposium on Research in Computer Security*. Springer, 716–736.

[7] Satadru Dey and Munmun Khanra. 2020. Cybersecurity of plug-in electric vehicles: Cyberattack detection during charging. *IEEE Transactions on Industrial Electronics* 68, 1 (2020), 478–487.

[8] Hossam ElHussini, Chadi Assi, Bassam Moussa, Ribal Atallah, and Ali Ghrayeb. 2021. A tale of two entities: Contextualizing the security of electric vehicle charging stations on the power grid. *ACM Transactions on Internet of Things* 2, 2 (2021), 1–21.

[9] Zacharenia Garofalaki, Dimitrios Kosmanos, Sotiris Moschoyiannis, Dimitrios Kallergis, and Christos Douligeris. 2022. Electric vehicle charging: A survey on the security issues and challenges of the open charge point protocol (OCPP). *IEEE Communications Surveys & Tutorials* 24, 3 (2022), 1504–1533.

[10] Christian Gorenflo, Ivan Rios, Lukasz Golab, and Srinivasan Keshav. 2017. Usage patterns of electric bicycles: an analysis of the WeBike project. *Journal of advanced transportation* 2017, 1 (2017), 3739505.

[11] Raju Gottumukkala, Rizwan Merchant, Adam Tauzin, Kaleb Leon, Andrew Roche, and Paul Darby. 2019. Cyber-physical system security of vehicle charging stations. In *2019 IEEE Green Technologies Conference (GreenTech)*. IEEE, 1–5.

[12] Erdem Gumrukcu, Ali Arsalan, Grace Muriithi, Charukeshi Joglekar, Ahmed Aboulebdeh, Mustafa Alparslan Zehir, Behnaz Papari, and Antonello Monti. 2022. Impact of cyber-attacks on EV charging coordination: The case of single point of failure. In *2022 4th Global Power, Energy and Communication Conference (GPECOM)*. IEEE, 506–511.

[13] Safa Hamdare, Omprakash Kaiwartya, Mohammad Aljaidi, Manish Jugran, Yue Cao, Sushil Kumar, Mufti Mahmud, David Brown, and Jaime Lloret. 2023. Cybersecurity risk analysis of electric vehicles charging stations. *Sensors* 23, 15 (2023), 6716.

[14] Md Nazmul Hasan, Rafia Nishat Toma, Abdullah-Al Nahid, MM Manjurul Islam, and Jong-Myon Kim. 2019. Electricity theft detection in smart grid systems: A CNN-LSTM based approach. *Energies* 12, 17 (2019), 3310.

[15] Paria Jokar, Nasim Arianpoo, and Victor CM Leung. 2015. Electricity theft detection in AMI using customers' consumption patterns. *IEEE Transactions on Smart Grid* 7, 1 (2015), 216–226.

[16] Dustin Kern, Christoph Krauß, and Matthias Hollick. 2023. Detection of anomalies in electric vehicle charging sessions. In *Proceedings of the 39th Annual Computer Security Applications Conference*. 298–309.

[17] Amirhossein Khazali, Yazan Al-Wreikat, Ewan J Fraser, Mobin Naderi, Matthew J Smith, Suleiman M Sharkh, Richard G Wills, Daniel T Gladwin, David A Stone, and Andrew J Cruden. 2024. Sizing a Renewable-Based Microgrid to Supply an Electric Vehicle Charging Station: A Design and Modelling Approach. *World Electric Vehicle Journal* 15, 8 (2024), 363.

[18] Tim Krause, Raphael Ernst, Benedikt Klaer, Immanuel Hacker, and Martin Henze. 2021. Cybersecurity in power grids: Challenges and opportunities. *Sensors* 21, 18 (2021), 6225.

[19] Ralph Langner. 2015. To Kill a Centrifuge-A Technical Analysis of What Stuxnet's Creators Tried to Achieve. 2013. *http://www. langner. com/en/wp-content/uploads/2013/11/To-kill-a-centrifuge. pdf>. Acesso em* 7 (2015), 21.

[20] Yang Liu, Yuchen Zhou, and Shiyan Hu. 2017. Combating coordinated pricing cyberattack and energy theft in smart home cyber-physical systems. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems* 37, 3 (2017), 573–586.

[21] Stephen McLaughlin, Brett Holbert, Ahmed Fawaz, Robin Berthier, and Saman Zonouz. 2013. A multi-sensor energy theft detection framework for advanced metering infrastructures. *IEEE journal on selected areas in communications* 31, 7 (2013), 1319–1330.

[22] Shanthan Kumar Padisala, Shashank Dhananjay Vyas, and Satadru Dey. 2024. Exploring Adversarial Threat Models in Cyber Physical Battery Systems. *arXiv preprint arXiv:2401.13801* (2024).

[23] Rouzbeh Razavi, Amin Gharipour, Martin Fleury, and Ikpe Justice Akpan. 2019. A practical feature-engineering framework for electricity theft detection in smart grids. *Applied energy* 238 (2019), 481–494.

[24] Gautam Rituraj, Gautham Ram Chandra Mouli, and Pavol Bauer. 2022. A comprehensive review on off-grid and hybrid charging systems for electric vehicles. *IEEE Open Journal of the Industrial Electronics Society* 3 (2022), 203–222.

[25] James ON Wilson and Tek Tjing Lie. 2022. Off-grid EV charging stations to reduce the impact of charging demand on the electricity grid. In *2022 7th IEEE Workshop on the Electronic Grid (eGRID)*. IEEE, 1–5.

[26] Yuqian Zhang, Fan Zhang, Yanjie Ji, and Yong Liu. 2023. Understanding the illegal charging intention of electric micro-mobility vehicle users by extending the theory of planned behavior. *Journal of Cleaner Production* 413 (2023), 137491.