

# A Privacy Preserving Approach to Energy Theft Detection in Smart Grids

Christopher Richardson and Nicholas Race  
School of Computing & Communications  
Lancaster University  
Lancaster, UK  
c.richardson@lancaster.ac.uk, n.race@lancaster.ac.uk

Paul Smith  
Digital Safety & Security Department  
AIT Austrian Institute of Technology  
Vienna, Austria  
paul.smith@ait.ac.at

**Abstract**—A major challenge for utilities is energy theft, wherein malicious actors steal energy for financial gain. One such form of theft in the smart grid is the fraudulent amplification of energy generation measurements from DERs, such as photovoltaics. It is important to detect this form of malicious activity, but in a way that ensures the privacy of customers. Not considering privacy aspects could result in a backlash from customers and a heavily curtailed deployment of services, for example. In this short paper, we present a novel privacy-preserving approach to the detection of manipulated DER generation measurements.

**Index Terms**—Smart grid; Smart metering; Energy theft; Detection; Privacy

## I. INTRODUCTION

Smart grids make use of Information and Communication Technology (ICT) and Supervisory Control and Data Acquisition (SCADA) systems to support the monitoring and automated control of Distributed Energy Resources (DERs) [1]. Alongside these sub-systems, an Advanced Metering Infrastructure (AMI) permits two-way communication between a smart meter at the customer premises and a utility company, such as a Distribution System Operator (DSO). Together, these systems can be used to collect data regarding energy usage and generation by so-called *prosumers* – customers that both consume and generate energy.

A major challenge for utilities is energy theft, wherein malicious actors steal energy for financial gain. Energy theft can have operational, safety and financial implications. There are many ways to commit energy theft [2], including altering the consumption and generation measurements that are collected from smart meters. In the former case, the attacker aims to under-report the amount of energy they consume (to reduce their bill); in the latter, the aim is to amplify the reported amount of energy generated, to increase remuneration from an energy supplier or aggregator, for example. It is necessary for utilities to deploy approaches to detect energy theft. However, it is important this is done in a way that respects the right to privacy of customers – not doing so could result in penalties and heavily curtailed deployment of services [3].

Previous work has investigated the prevention and detection of attacks to the AMI [4]–[8]. For example, Jokar and Leung have developed a pattern-based energy theft detector, which compares the output from transformer meters for a neighbourhood with values from smart meters [6]. They train a

Support Vector Machine (SVM) with historical data for each customer. Meanwhile, Faisal *et al.* [7] propose an intrusion detection system that makes use of multiple points in an AMI architecture to detect attacks using a number of classifier algorithms. Liu *et al.* [8] propose a means of detecting data injection attacks to smart meters, which operates on the meter itself. Arguably, the cost of adding such functionality to smart meters makes it unlikely to be used in reality. Importantly, the authors of these contributions do not consider privacy aspects.

To address privacy concerns in the smart grid [9], a number of approaches have been proposed [10]–[12]. For example, Chen *et al.* advocate a scheme to modulate the behaviour of a water heater to give the impression someone is always home [11]. They claim their approach does not waste energy, but simply modifies when water is heated. Such a scheme could prove problematic in the presence of demand-response services, for example. Efthymiou and Kalogridis [12] propose a mechanism for smart meter reading that allows the secure attribution of measurements to a location, not a specific customer. This allows utilities to use meter readings for operational purposes. In our case, we require the capacity to identify customers that may be acting nefariously, but do not require (or rather, do not want to know) information about their private usage patterns.

In this short paper, we present a novel privacy-preserving approach to the detection of manipulated DER power generation measurements. Our aim is to detect malicious actors that over-report the power they generate for financial gain, in a way that ensures private information [13], which is contained in smart meter measurement data, is not revealed to third-parties. Our approach leverages the insight that (normalised) power output from photovoltaic (PV) installations in a geographical region should be similar – deviations from the norm could indicate malicious behaviour. In short, our approach calculates the Euclidean distance between all pairs of normalised power generation measurements from PV installations in a region. Subsequently, a clustering algorithm is used to identify outlying distances that could indicate malicious behaviour. Privacy is ensured as only the Euclidean distances and homomorphically encrypted measurements need to be shared with third-parties.

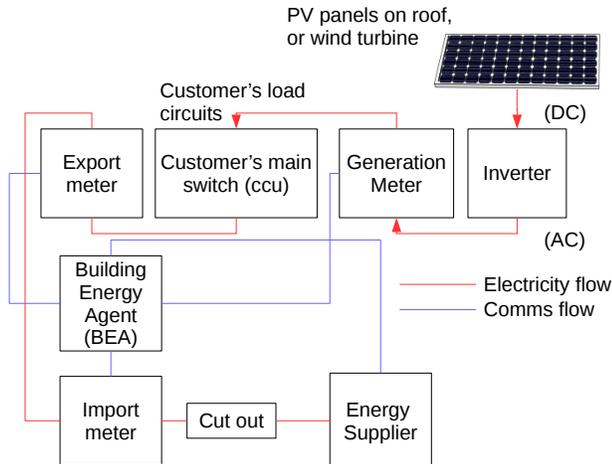


Fig. 1. Schematic of a typical residential photovoltaic installation, showing the electricity and communication connections [14]

## II. A SMART GRID ENERGY THEFT SCENARIO

In this section, we present the smart grid deployment scenario that we are considering and the attacker model that we assume for our work.

### A. Smart Grid Deployment

The scenario we consider includes numerous residential properties that have Distributed Energy Resources (DERs) installed, such as photovoltaics (PVs). So-called *prosumers* both consume and produce energy. Figure 1 presents a typical deployment of a PV installation at a residential property [14], showing electrical and communication connections. In this setting, the PV panels produce DC power, which is converted to AC power by an Inverter. A Generation Meter records the energy generated by the PV installation; this meter may be integral to the Inverter. The output of the Generation Meter goes to the Customer Consumer Unit (CCU). The Export Meter records the energy that is fed into the electricity grid, whilst the Import Meter records what is being consumed. Data from both the Export Meter and Generation Meter are sent to a prosumer's energy supplier for billing and remuneration purposes<sup>1</sup>. The Building Energy Agent (BEA) runs energy services, e.g., it participates in demand-response services and manages the generation and controllable loads at the premises. Typically, this device is connected to a wide-area network, such as the Internet, to support the implementation of these services.

### B. Attacker Model

McLaughlin *et al.* [16] have developed an attacker model that describes the types of attackers who may be sufficiently

<sup>1</sup>For example, E.ON, a UK energy supplier, states that prosumers will receive payment for every kWh that is generated, whether it is used locally or exported to the grid [15].

motivated to commit energy fraud – they highlight that customers, e.g., prosumers, may be inclined to reduce their energy bills in this way. In addition to examining threat sources and their motivation, they consider the ways in which smart meters could be attacked. For example, they have identified that data can be manipulated in transit between the meter and the supplier or when it is stored on the smart meter. Additionally, by tampering with the smart meter, specifically the Export or Generation Meter, it is possible for a prosumer to manipulate the data that is communicated to the energy supplier, in order to fraudulently claim they are producing more energy than they are in actuality.

A way to detect such attacks involves identifying anomalous measurements, out of a larger set, which could indicate malicious behaviour. However, to use the data in this way, there are a number of privacy issues that must be addressed. For example, if the PV output is known along with the amount of energy that is fed into the grid, it may be possible to determine when a prosumer is at home, as the feed-in energy will likely be lower if they are using appliances (assuming they are using the power they generate). Consequently, technologies that ensure the privacy of prosumers must be implemented, alongside techniques that can be used to identify malicious behaviour.

## III. AN APPROACH TO PRIVACY-PRESERVING ENERGY THEFT DETECTION

In this section, we describe an approach to detecting the fraudulent manipulation of electricity generation data, and a scheme that builds on this approach to ensure privacy.

### A. Energy Theft Detection Approach

The insight that we use as a basis for detecting malicious activity is that the behaviour of PV installations, in terms of their normalised energy output, is geographically correlated. To support this claim, we performed an analysis of a dataset that was collected from residential PV installations in the UK<sup>2</sup>. For each installation, the dataset contains periodic energy generation measurements (kWh), which are recorded every thirty minutes, along with the latitude, longitude, and its peak generation capacity (kWp). To calculate the normalised energy output  $P$  for a given time period  $t_i$ , we divide the kWh measure for the period by the installation's peak capacity (see Eq. 1).

$$P_{t_i} = \frac{kWh_{t_i}}{kWp} \quad (1)$$

For this analysis, we segmented the entire region under evaluation into a 4x4 grid. Figure 2 depicts how the PV installations in the dataset are geographically distributed. The colour of a point in the figure indicates the normalised energy output at midday. Figure 3 shows the normalised measurements from these PV installations for a single day, with each plot representing a distinct geospatial region. Each

<sup>2</sup>The data was obtained from the Sheffield Solar Group at the University of Sheffield: <http://www.microgen-database.org.uk/>

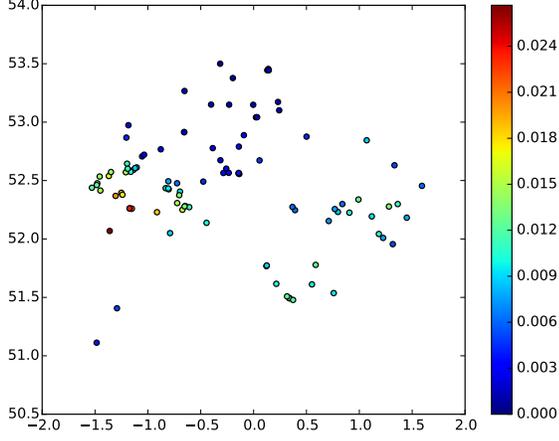


Fig. 2. The geospatial distribution of the photovoltaic installations in the dataset we used, showing normalised energy output at midday for a selected day; the x- and y-axis and longitude and latitude values.

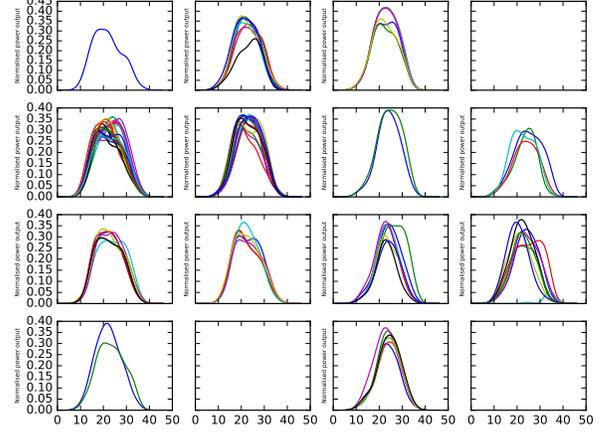


Fig. 3. Normalised energy output for a day for PV installations in a geospatial area. The x-axis shows measurements over time; the y-axis are normalised energy measurements. Empty plots contain no PV installations.

curve represents the normalised energy output from a single installation. It can be seen from this analysis that PV behaviour appears to be geographically correlated. In other data, not shown here for space reasons, we have seen these normalised values to be regionally distinct. We can use this as a basis for identifying measurements that deviate from the norm.

In order to detect anomalous measurements, it is first necessary to segment the PV installations into geospatial regions, allowing us to then compare the power output from those in a similar vicinity. After performing this segmentation, we calculate the Euclidean distance between the normalised PV power output of two installations over a single day; this is performed for all combinations of installation pairs in the geographic region. The Euclidean distance is calculated using Eq. 2, in which  $P_1 t_i$  and  $P_2 t_i$  are the normalised energy output measurements for two installations for time period  $t_i$ , and  $n$  is the total number of measurements that are taken over a day.

$$d(P_1, P_2) = \sqrt{\sum_{i=0}^n (P_{1t_i} - P_{2t_i})^2} \quad (2)$$

Subsequently, we cluster the Euclidean distance measurements in a region using the DBSCAN clustering algorithm [17], in order to identify outliers in a specific geospatial region. Outliers could be indicative of malicious behaviour. To increase the certainty that a particular installation is behaving in an anomalous way (and is therefore potentially malicious), we perform the analysis across multiple days.

### B. Ensuring Prosumer Privacy

In the UK, energy customers, such as prosumers, can choose the frequency that smart meter data is collected for billing purposes. The options that are available include monthly (a mandatory minimum), daily, or every thirty minutes. For privacy reasons, a customer may choose to have their meter

data read, as a cumulative total, on a daily or monthly basis. Meter data collection at these frequencies cannot be readily used to detect fraudulent behaviour using a scheme such as the one proposed in this paper.

To address this issue, and to ensure measurement data that could reveal private information is not sent to third-parties, we have developed a privacy-preserving means of implementing our detection approach. The technique results in just a Euclidean distance measure being sent to an energy supplier (or other suitable third-party), which is used to detect outliers and malicious behaviour. These distances cannot be used to infer detailed energy consumption or generation profiles. The Euclidean distances are computed by Building Energy Agents (BEAs) in the field, using homomorphic computation. This combination ensures that generation metering data, which could reveal personally sensitive information, are not distributed in a way that can be openly read by third-parties. However, in this way, the detection of malicious behaviour can still be achieved.

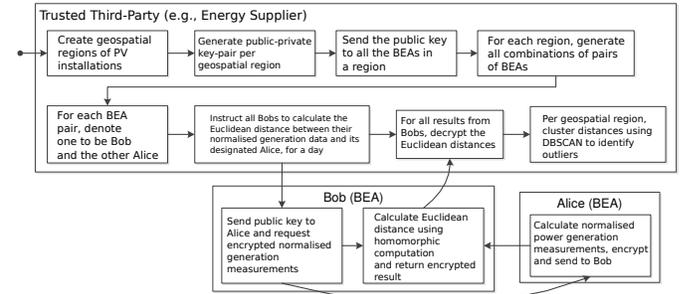


Fig. 4. Euclidean distance calculation with Paillier cryptosystem

Figure 4 depicts how our privacy scheme works. Initially, the energy supplier creates a public-private key using the Paillier cryptosystem [18]. The public key is sent to all

the BEAs that reside in a pre-calculated geospatial region of PV installations. (The public key is used to encrypt the normalised energy and Euclidean distance measurements.) It is then necessary for the Euclidean distance to be calculated for all combinations of installation pairs. To accomplish this, a list of PV installation pairs – called Alice and Bob for explanatory purposes – are generated by the energy supplier. Subsequently, each Bob in a pair is instructed to calculate the Euclidean distance between its own measurements and those from Alice, for the full day, using Eq. 2. To achieve this, Alice sends its normalised energy generation measurements as ciphertext to Bob. Using a method for homomorphic computation, proposed by Rane *et al.* [19], Bob is able to calculate the Euclidean distance between the measurements. Once the resulting ciphertext has been generated by Bob, it is sent to the energy supplier. Finally, the supplier decrypts the result from Bob, with its private key, and uses a clustering algorithm to detect outlying distances within a region.

#### IV. INITIAL EVALUATION

To evaluate the effectiveness of our detection approach, we conducted experiments using the UK dataset. We created ten geospatial regions, containing on average ten installations. We introduced a single malicious installation into each region, which over-reported their generation measurements. We executed our detection algorithm over a year’s worth of data, using the same malicious actor for each day. The F-score was calculated for the complete year’s worth of results from a single day. When a malicious actor reports twice their actual generation, we see an F-score of 0.9, whereas for three times the value is 0.97 – an ideal F-score is 1. These initial results indicate the detection approach can effectively identify this form of energy theft.

#### V. CONCLUSION

A major challenge for utilities is energy theft, wherein malicious actors steal energy for financial gain. One such form of theft in the smart grid is the fraudulent amplification of energy generation measurements from DERs, such as photovoltaics.

In this short paper, we have introduced a novel privacy-preserving approach to detecting such attacks. It builds on the observation that normalised energy output from PV installations in a geographic region are similar. Malicious behaviour is detected by calculating the Euclidean distance between energy output measurements from an installation over a day. These distances are then clustered to identify outliers and potentially malicious behaviour. Privacy is preserved through the use of homomorphic computation to allow Euclidean distance measures to be calculated privately – this distance is the only information that is given to a third-party, such as an energy supplier.

In our ongoing work, we are evaluating two major aspects of our approach: (i) the overhead associated with performing calculations in a distributed manner on BEAs; and (ii) the detection performance of our approach. Initial results indicate the approach has promise. Future work will investigate whether our scheme could be applied to other forms of DER,

such as wind turbines, and whether faults in installations can be identified, as well as malicious behaviour.

#### ACKNOWLEDGMENTS

This research has been supported by funding from the European Union Seventh Framework Programme (FP7/2007-2013) under grant agreement No. 608224.

#### REFERENCES

- [1] Y. Yan *et al.*, “A survey on smart grid communication infrastructures: Motivations, requirements and challenges,” *IEEE Communications Surveys Tutorials*, vol. 15, no. 1, pp. 5–20, January–March 2013.
- [2] R. Czechowski and A. M. Kosek, “The Most Frequent Energy Theft Techniques and Hazards in Present Power Energy Consumption,” in *Joint Workshop on Cyber-Physical Security and Resilience in Smart Grids (CPSR-SG 2016)*, Vienna, Austria, April 2016.
- [3] R. Hoenkamp *et al.*, “The Neglected Consumer: The Case of the Smart Meter Rollout in the Netherlands,” *Renewable Energy Law and Policy (RELPL)*, no. 4, pp. 269–282, November 2011.
- [4] R. Berthier and W. H. Sanders, “Specification-Based Intrusion Detection for Advanced Metering Infrastructures,” in *17th IEEE Pacific Rim International Symposium on Dependable Computing (PRDC)*, December 2011, pp. 184–193.
- [5] M. LeMay *et al.*, “Unified Architecture for Large-Scale Attested Metering,” in *40th Annual Hawaii International Conference on System Sciences (HICSS 2007)*, January 2007, p. 115.
- [6] P. Jokar, N. Arianpoo, and V. C. Leung, “Electricity theft detection in ami using customers consumption patterns,” *IEEE Transactions on Smart Grid*, vol. 7, no. 1, pp. 216–226, 2016.
- [7] M. A. Faisal *et al.*, “Securing advanced metering infrastructure using intrusion detection system with data stream mining,” in *Intelligence and Security Informatics Pacific Asia Workshop (PAISI 2012)*, Kuala Lumpur, Malaysia, May 2012, pp. 96–111.
- [8] X. Liu *et al.*, “A Collaborative Intrusion Detection Mechanism Against False Data Injection Attack in Advanced Metering Infrastructure,” *IEEE Transactions on Smart Grid*, vol. 6, no. 5, pp. 2435–2443, September 2015.
- [9] A. Molina-Markham *et al.*, “Private memoirs of a smart meter,” in *2nd ACM workshop on Embedded Sensing Systems for Energy-efficiency in Buildings*. ACM, 2010, pp. 61–66.
- [10] S. Finster and I. Baumgart, “Privacy-Aware Smart Metering: A Survey,” *IEEE Communications Surveys Tutorials*, vol. 16, no. 3, pp. 1732–1745, May 2014.
- [11] D. Chen *et al.*, “Preventing occupancy detection from smart meters,” *IEEE Transactions on Smart Grid*, vol. 6, no. 5, pp. 2426–2434, Sept 2015.
- [12] C. Efthymiou and G. Kalogridis, “Smart Grid Privacy via Anonymization of Smart Metering Data,” in *2010 First IEEE International Conference on Smart Grid Communications (SmartGridComm)*, October 2010, pp. 238–243.
- [13] U. Greveler *et al.*, “Multimedia content identification through smart meter power usage profiles,” *Computers, Privacy and Data Protection*, 2012.
- [14] T. Chevalier, in *Operation of electricity meters when energy flowing in reverse*. Association of Meter Operators, 2013.
- [15] E.ON, “Feed-in tariff factsheet,” [https://www.eonenergy.com/media/PDFs/For-your-home/Feed-in-Tariffs/Feed\\_In\\_fact\\_sheet\\_0172\\_04\\_12\\_Web.pdf](https://www.eonenergy.com/media/PDFs/For-your-home/Feed-in-Tariffs/Feed_In_fact_sheet_0172_04_12_Web.pdf).
- [16] S. McLaughlin *et al.*, “Energy theft in the advanced metering infrastructure,” in *Critical Information Infrastructures Security*. Springer, 2010, pp. 176–187.
- [17] D. Birant and A. Kut, “St-dbscan: An algorithm for clustering spatial-temporal data,” *Data & Knowledge Engineering*, vol. 60, no. 1, pp. 208–221, 2007.
- [18] P. Paillier, “Public-key cryptosystems based on composite degree residuosity classes,” in *Advances in cryptography - EUROCRYPT99*. Springer, 1999, pp. 223–238.
- [19] S. D. Rane *et al.*, “Secure distortion computation among untrusting parties using homomorphic encryption,” in *Image Processing (ICIP), 2009 16th IEEE International Conference on*. IEEE, 2009, pp. 1485–1488.