Pacific Northwest
NATIONAL LABORATORY

PNNL-35257

A review of privacy in energy applications

October 2023

Sebastian Cardenas, D. Jonathan
Mukherjee, Monish
Ramirez, Javier

U.S. DEPARTMENT OF
ENERGY

Prepared for the U.S. Department of Energy
under Contract DE-AC05-76RL01830

DISCLAIMER

This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor Battelle Memorial Institute, nor any of their employees, makes **any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights**. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof, or Battelle Memorial Institute. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

PACIFIC NORTHWEST NATIONAL LABORATORY
operated by
BATTELLE
for the
UNITED STATES DEPARTMENT OF ENERGY
under Contract DE-AC05-76RL01830

Printed in the United States of America

Available to DOE and DOE contractors from
the Office of Scientific and Technical Information,
P.O. Box 62, Oak Ridge, TN 37831-0062
www.osti.gov
ph: (865) 576-8401
fox: (865) 576-5728
email: reports@osti.gov

Available to the public from the National Technical Information Service
5301 Shawnee Rd., Alexandria, VA 22312
ph: (800) 553-NTIS (6847)
or (703) 605-6000
email: info@ntis.gov
Online ordering: http://www.ntis.gov

PNNL-35257

A review of privacy in energy applications

October 2023

Sebastian Cardenas, D. Jonathan
Mukherjee, Monish
Ramirez, Javier

Prepared for
the U.S. Department of Energy
under Contract DE-AC05-76RL01830

Pacific Northwest National Laboratory
Richland, Washington 99354

Abstract

As the distribution system continues to experience an increase in distributed energy resource (DER) and electric vehicle (EV) penetration, so does the need for new solutions that can help grid operators manage and leverage their capabilities. This will undoubtedly lead to new operational schemes and business opportunities that will transform the traditional consumer into a prosumer who will be more actively engaged in grid operations. Although the field is still under active development, many of the potential use cases presented in literature or industry are built upon edge computing, two-way communications, and other innovative computational constructs to attain their goals. However, at their core, many use cases assume a great level of data access to aid with the decision-making process, an assumption that may need to be revised to ensure fair and equitable operational processes are maintained. This may be particularly true as edge resources are predicted to participate in retail-side, many-to-many, or peer-to-peer markets and thus may lead to financial impacts if data access considerations are ignored.

The need to revise data access mechanisms can be further justified by the introduction of new participants into the operational process, who do not have the same level of trust, nor the incentives to focus on energy delivery as their primary objective. At the same time, more consumers are becoming aware of their own data, and the potential impacts of its abuse. To help solution developers better understand these risks, this report has been developed to offer an initial introduction to the topic of privacy. This is achieved by 1) Highlighting the need for privacy-aware solutions; 2) Encouraging system designers to be inquisitive about the status quo; 3) Documenting the existing threat space; 4) Presenting and evaluating tools that may be helpful towards enabling better privacy postures; and 5) Making recommendations to encourage the adoption of better practices.

From a technical perspective, the report focuses on evaluating two potential techniques by applying them to the Transactive Energy space. Based on the obtained results, it can be established that Differential Privacy (DP) methods may have limited applicability when highly correlated, time-series data records need to be protected. However, DP may be a powerful tool when it is used to aggregate and analyze mid-size and large-size data sets in a more traditional statistical environment. The second tool under evaluation is threshold cryptography, which can guarantee complete secrecy (and thus privacy) but requires the establishment of key management procedures and dedicated communication channels for key coordination. Therefore, due to its increased computational overhead, the use of threshold cryptography must be weighted using a cost/benefit analysis on a per-application basis.

Summary

The successful integration of renewables will require engineers to develop new communication architectures and processes to effectively harness the capabilities provided by DERs, EVs, and other emerging technologies. However, these new processes need to be mindful of the different and sometimes contradictory requirements of each actor, thereby helping to foster fair and equitable business practices. One of the emerging aspects is digital privacy, which is a set of expectations that an individual has over their digital footprint, which usually covers the data and associated metadata generated during day-to-day business transactions. As expected, multiple researchers have proposed a variety of mechanisms designed to protect the data across its operational lifecycle, thereby helping to increase privacy guarantees for data producers.

One of the promising technologies is Differential Privacy (DP), which offers a mathematical framework that enables individuals to attain strong privacy guarantees. Such guarantees are intended to encourage participants to share data more openly since their inputs within a given population are guaranteed to be truly anonymous. However, the method assumes that individuals participating in such groups are independent from other members, a requirement that may be hard to satisfy in certain grid applications. To demonstrate this limitation, this work presents a quantitative evaluation of Differential Privacy when it is used to protect aggregated demand records. The results indicate that DP may underperform at protecting time-series data in the long-term but may be a viable option if data is collected over finite time periods.

In addition, this work also presents high-level recommendations to encourage system developers to better understand the limitations and issues associated with existing privacy practices and recognize the potential benefits of moving towards a privacy-by-design model that encourages secure data sharing. By adopting secure privacy principles, actors could increase their willingness to participate in new or novel business use cases that could help address future grid challenges.

Acronyms and Abbreviations

AMI	Advanced Metering Infrastructure
CNN	Convolutional Neural Network
DER	Distributed Energy Resource
DKG	Distributed Key Generation
DLT	Distributed Ledger Technologies
DP	Differential Privacy
EULR	End Use Load Research
EV	Electric Vehicle
GSP	Graph Signal Processing
HMM	Hidden Markov Model
HVAC	Heating, Ventilation, and Air Conditioning
PV	Photovoltaic
TES	Transactive Energy Systems
TC	Threshold Cryptography
TOU	Time of Use
ZKP	Zero-Knowledge Proof

Acknowledgments

This project was supported by the Department of Energy, Office of Electricity, Advanced Grid Research and Development Program. The authors would like to thank Chris Irwin for his support and contributions to shaping the scope and direction of this work.

Contents

Figures

Tables

1 Introduction

For most of the electric industry’s history, the utility customer was seen as the final destination for energy delivery. From a customer’s perspective, the local utility had the responsibility to deliver energy and to produce invoices based on consumption records. Before the deployment of automatic meter reading technology and the introduction of Time of Use (TOU) tariffs, most utilities collected a finite set of consumption records during the year (e.g., monthly meter readings), which provided limited insight into consumer behaviors. However, as measurements started to be collected more often, utilities started to gain finer visibility into their customer habits, enabling them to optimize operational processes and offer enhanced feedback to their customers.

Although technologies, such as Advanced Metering Infrastructure (AMI) are known for their capability to increase visibility into the distribution system and help support new or enhanced use cases, some customers have expressed objections. Some objections, as expressed by customers include: uncertainty in monthly payments due to new or alternative rate structures, unwarranted direct control actions during demand response events, and risk related to the invasion of privacy (Darby, Sarah J, 2012) (Pepermans, 2014). Some of the privacy concerns have been addressed by tweaking and leveraging existing regulatory frameworks that utilities must abide by. Although most regulations are purely based on technical needs (e.g., for being part of critical infrastructure) some may be guided by social policies (e.g., limiting rate hikes). For example, in the Netherlands, in response to public objections utilities may only access cumulative metering to compute energy costs even if higher resolution readings are technically available (Darby, Sarah J, 2012). In contrast, Ontario (Canada) has achieved a balance between use and privacy which has been attributed to the local utilities’ reach-out programs and their previously earned trust that enabled social acceptance with regards to AMI deployments.

Within the USA, privacy requirements are set by state-level energy commissions, leading to varying levels of consumer-level privacy protections and legal enforcement capabilities (Dasom & Hess, 2021). While states with a strong history of consumer-oriented privacy guidelines may have comprehensive laws that limit the sharing of consumer-level data to third parties without authorization (e.g., California), other states do not have any laws nor proposed legislative bills to address the topic. This uneven landscape may risk hampering the grid digitalization process by undermining social acceptance of new data-driven operational and business models if consumers cannot be guaranteed adequate levels of protection.

Although from a policy perspective, there remains a long way towards achieving effective privacy protections, recent attention to the topic could enable data producers to gain a greater level of awareness over their digital footprint. Generally speaking, data-driven grid participants may ask themselves the following questions:

Where can data be obtained? This may define the mechanisms and procedures that must be followed before access is granted. This may include the protection mechanisms that need to be applied to the data records (e.g., adding digital watermarks to track and prevent re-sharing).

When can data be accessed? In general, from a business perspective, up-to-date data has a higher value than stale data. As such, consumers can define data embargo periods that balance their own privacy needs with their economic value. Examples may include, granting access after 3 months or limiting access to only 12 months of data.

What data can be viewed? Data producers may choose the types of data they're willing to share (instantaneous, cumulative), their granularity (1 minute, 5 minutes, 1 day, etc.), their discretization levels (W, kW, or by classification bin), and its accuracy (amount of noise).

Who can access data? As consumers start to engage with more parties, they may need to implement selective access mechanisms to accommodate the different actors and their needs. Data examples may include requiring interested parties to justify their need for access, requiring third parties to notify data producers when data breaches are identified, and defining boundaries on the re-distribution of raw or processed datasets (similar to copyleft restrictions in software licenses).

Why is data being produced? A re-evaluation of the actual needs versus the implemented or planned deployment must be carried out to ensure data being produced is driven by actual needs. This may include down-selecting and filtering individual telemetry streams, down-sampling high-frequency streams, or randomizing the sampling interval to further protect end-user privacy.

Clearly, some of the aforementioned considerations require a system-level perspective to balance all the needs and wants of all relevant actors. This may include re-thinking existing processes and de-facto operational modes to ensure data being produced, transported, and consumed is in fact driven by needs. Although radical changes to existing processes may be cost-prohibitive, grid transformation opportunities, such as those driven by the ongoing vehicle electrification process and the emergence of regulations that enable DERs to provide grid services should be seized as an ideal opportunity to build and deploy privacy-by-design solutions. These privacy-by-design solutions will require solution developers to satisfy business functionalities while making sure they remain compliant with the end-user[1] protections provided by the legislative framework as well as other customer-defined privacy preferences.

Based on this assumption, this report focuses on describing and evaluating privacy techniques that can be used to construct privacy-by-design software solutions within the energy space. The presented techniques can be broadly categorized into:

- **Mathematical solutions:** Methods that seek to prevent individual data records from being extracted given a collection of aggregated records. This allows data consumers to make operational decisions while protecting the privacy of the data producer(s). It is important to note that aggregated records may involve a) An individual's data record that has been combined with other individuals' records to create a group, or b) A set of records captured at distinct times to represent an individual's behavior during a given timespan.

- **Cryptographic solutions:** Methods that can ensure a record's confidentiality until the data producer (or other authorized entity) releases the encryption key to the data consumer. Depending on the method being used, two levels of confidentiality may be achieved: a) complete confidentiality – *the record cannot be distinguished from a random bit sequence* or b) partial confidentiality – *certain metadata or high-level operations can be carried out without revealing the original record.*

[1] It is important to note that in this report the term end-user is used to describe any physical or legal entity that produces or consumes data records. This for example may include a retail consumer/prosumer, or a business partner who shares operational records with the local grid operator.

The report is organized as follows, Section 2 provides an overview of the existing privacy landscape. It introduces key risks that are relevant to the energy sector along with potential solutions. Section 3 focuses on describing and evaluating two potential privacy constructs that were selected for further analyses, while Section 4 summarizes the recorded observations. Finally, Section 5 is used to capture conclusions as well as provide future work ideas.

2 State of the Art

The topic of digital privacy is often referenced in the context of Internet-driven business and other tech-related solutions. This may include web-based trackers that can be used to provide personalized ads, location trackers that can be used to analyze mobility patterns, and service aggregators that can link multiple data sources to assemble complex profiles that may enable linking digital accounts to physical identities. However, with the increasing digitalization of the energy sector and the emergence of new business models, resulting data exchanges could be collected and analyzed to reveal cyber-physical interactions. This may include determining the presence and usage patterns of customer-level assets, including EV usage or Photovoltaic (PV) production values, and, in general, observe how different actors[1] respond to a wide array of events or signals (e.g., meteorological events, or macro-economic signals).

In the current landscape, most customers who are part of a traditional market structure have a strong dependency on the local grid operator, who has the responsibility to ensure a continuous supply of electricity in exchange for a service fee. As expected, most utility revenue generation mechanisms focus on ensuring full cost recovery, reducing operational costs, and minimizing energy procurement costs. In addition, due to local regulations that discourage price discrimination, direct monetization of consumer-originated data has been limited, other than being used to improve forecasts or support other cost optimization solutions. However, as local area power markets start to emerge, new actors with diverse financial incentives will gain access to grid data, potentially leading to the unexpected exploitation of private data.

Although the privacy risks will ultimately depend on the underlying use case and the actual implementation, sample scenarios that could lead to abuse include:

- Combining multiple data sources to de-anonymize data producers, thereby enabling sophisticated actors to correlate seemingly disparate activities or records and tie them back to digital or physical identifiers. This may allow third-party actors to link consumption behaviors to individual service addresses or customer identities.

- Using data analytics in a microgrid or peer-to-peer energy market to identify a competitor’s business strategy and other cost-response relations to predict market behaviors. Systematic analysis of competitors’ behavior could lead to unfair market dominance as well as cause other disruptions that reduce participants’ confidence.

Similarly, well-designed privacy-aware solutions could be used to:

- Facilitate the adoption of bi-directional communication processes to help support service negotiations between two or more system actors at the edge. Privacy-aware signal exchanges can allow participants to assemble and exchange grid services while preventing external agents from obtaining private information.

- Enable higher-level supervisory actors to attain system visibility without requiring raw data access, thereby helping end-users to create a positive impact while minimizing personal risks.

[1] Actors in this context may involve customers, prosumers, service brokers, grid operators, and other relevant stakeholders who share data for operational reasons.

- By adopting privacy-aware collection procedures during the data procurement lifecycle, organizations may be able to limit the potential impacts associated with data leaks in case the records are compromised in the future.

Clearly, privacy is a desirable system property that, if well implemented, could be used as a tool to enable new operational schemes. Nevertheless, the successful adoption of privacy protection schemes requires engineers to evaluate each of the underlying construct features and characteristics to ensure application's needs and expectations are being met. This may be particularly challenging to achieve in the energy sector, which requires solutions to remain robust, and scalable while being able to operate over long-periods of time (i.e., decades instead of the typical 3 to 5 years market cycle seen in the Information Technology domain). To aid with this topic, this section presents a high-level overview of the available data anonymization constructs, some of their common pitfalls, as well as promising technologies that could be used within the energy sector.

2.1 Anonymization, An Overview

Data anonymization is the process of altering or removing personal or otherwise sensitive information from a dataset to prevent the identification of individuals. The goal is to transform the data in such a way that it becomes difficult or impossible to associate specific traits or records with a particular person while still maintaining the overall utility of the dataset for analysis or other purposes . Anonymization may include the permanent modification or deletion of data with the intent of preventing non-authorized actors to reverse the mapping function. Solutions may use a combination of suppression, generalization, masking, and other related techniques to achieve desired goals. A brief overview of these techniques can be summarized as follows:

- **Suppression:** The technique seeks to remove uniquely identifiable traits from the record itself. Examples may include removing the customer's name or service address from the data point itself. It may not be well suited for data consumers who require some level of data attribution (e.g., to track the validity of the source).
- **Generalization:** Sensitive fields are replaced with generalized values that enable data consumers to classify and group information without identifying individualistic attributes. Examples may include changing the service address to a region-based label, or an individual's name with a population *class* that captures their life stage, income bracket, or any other desired attribute.
- **Distortion:** A controlled amount of noise is added to the data records to provide a layer of anonymity. The amount of added noise may be provided by a one-way function (i.e., the original data point cannot be recovered), or via a two-way function (i.e., an intermediary can add or remove noise to enable data consumers to re-identify the data producers)
- **Swapping:** When multiple data producers are available, it becomes possible to swap individual data points among group members without affecting group-level analysis results. Although individual records may be protected, the group's privacy remains unprotected.
- **Masking:** As its name implies, the data masking process substitutes real data with fictitious data that enables data consumers to link records to a fake, but consistent identifier.

As expected, each technique has a set of potential benefits and drawbacks that must be analyzed under the context of the target application. In order to achieve an adequate level of protection, system designers may need to combine two or more core techniques to achieve the desired data safety level. Reliance on a single technique may lead to insecure implementations, as demonstrated by successful deanonymization attacks (Narayanan & Shmatikov, 2008) as well as other self-recognizant works that have identified deficiencies in past implementations (Long, 2020), (Hawes, 2021).

2.2 Attacks On Anonymization Efforts

Although re-mapping an individual's anonymized record to a unique identifier is often considered as a prime example of a deanonymization attack, other partial mappings such as, "an individual is a member of group m" or "individuals are members of group m" can result in privacy violations that may lead to unfair or undesirable system behaviors. Re-mapping is often possible when feature-rich datasets contain fields that are distinct enough to be considered unique on their own, or when unique, compounded keys can be assembled from subfields (e.g., by combining orthogonal fields such as age and zip code information). A review of re-identification attacks within the medical field has been documented in (El Emam, Jonker, Arbuckle, & Malin, 2011). The authors concluded that a significant number of records have been successfully re-identified by researchers (>25%), however, they also note that the reported percentages are inherently biased and should not be used to make generalized assumptions. Sources of bias may include research studies that have not been published due to poor re-identification rates, but also successful attacks that have remained unreported (e.g., due to ethical considerations, or because of their non-academic nature).

Although re-identification is often measured under a pass-or-fail grading scheme, partial or approximate re-identifications can also lead to problematic outcomes. Such scenarios can arise when data consumers can infer individual attributes with a low margin of error, or when small subsets of the population can be tagged with highly specific attributes. Such issues could allow data consumers to make highly informed guesses about the underlying population that could effectively circumvent a system's privacy guarantees. The use of indirect inference has for example been presented in (Lindholm, Richman, & M.V., 2022) to illustrate how discriminatory pricing can be applied to insurance premiums by relying on non-discriminatory attributes (i.e., zip code data). Based on this result, it could be inferred that vendors in the energy space could use income-based attributes to drive product offerings based on their clients' ability to pay while appearing to rely on technical requirements to artificially justify their needs.

Common attack tactics include:

Linkage Attack: A linkage attack is a method used by an attacker to link publicly available information with entries from an anonymized data set (Oak Ridge National Laboratory, 2023). Two prime examples of linkage attacks are homogeneity and external knowledge attacks (di Vimercati, 2023). A homogeneity attack can be performed when compound keys can be generated and cross-matched across two distinct datasets, further increasing the richness of the data (see Figure 1, for a graphical example). An external knowledge attack is performed when an attacker uses known information to reduce the population space, potentially enabling it to identify a sensitive characteristic in a probabilistic manner. For example, if consumers with EVs need to be identified, one may focus on higher-income individuals.

Membership Inference Attack: An attack on machine learning algorithms that leverages a model's tendency to perform better on data that they have already seen. Hence, the attacker's goal is to determine if a given data record entry is a member of the dataset originally used to train the model (Shokri, 2017). The attack is based on developing shadow models that have statistical properties similar to those used in training. Such a process may involve the development of adversarial machine learning models whose goal is to generate inputs that are representative of the original dataset.

Query Attack: Such an attack occurs when an attacker queries a protected dataset with the goal of extracting less-protected subsets that can later be assembled to reveal previously protected data. Further refinements of these methods have led to techniques such as the inversion attack, where the attacker constantly issues queries to a machine learning model in order learn the relationships between the output and the input (Wang, 2021), potentially helping to infer private information about the participants.

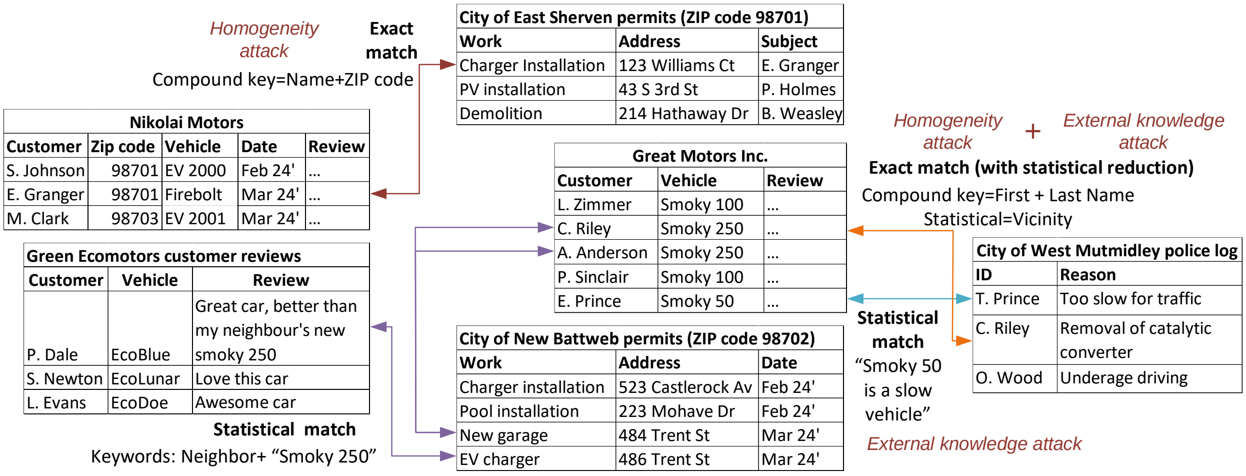
Homogeneity attack
Exact match
Compound key=Name+ZIP code
Homogeneity attack + External knowledge attack
Exact match (with statistical reduction)
Compound key=First + Last Name
Statistical=Vicinity
Statistical match
"Smoky 50 is a slow vehicle"
External knowledge attack
Statistical match
Keywords: Neighbor+ "Smoky 250"

City of East Sherven permits (ZIP code 98701)		
Work	Address	Subject
Charger Installation	123 Williams Ct	E. Granger
PV installation	43 S 3rd St	P. Holmes
Demolition	214 Hathaway Dr	B. Weasley

Nikolai Motors				
Customer	Zip code	Vehicle	Date	Review
S. Johnson	98701	EV 2000	Feb 24'	...
E. Granger	98701	Firebolt	Mar 24'	...
M. Clark	98703	EV 2001	Mar 24'	...

Green Ecomotors customer reviews		
Customer	Vehicle	Review
P. Dale	EcoBlue	Great car, better than my neighbour's new smoky 250
S. Newton	EcoLunar	Love this car
L. Evans	EcoDoe	Awesome car

Great Motors Inc.		
Customer	Vehicle	Review
L. Zimmer	Smoky 100	...
C. Riley	Smoky 250	...
A. Anderson	Smoky 250	...
P. Sinclair	Smoky 100	...
E. Prince	Smoky 50	...

City of New Battweb permits (ZIP code 98702)		
Work	Address	Date
Charger installation	523 Castlerock Av	Feb 24'
Pool installation	223 Mohave Dr	Feb 24'
New garage	484 Trent St	Mar 24'
EV charger	486 Trent St	Mar 24'

City of West Mutmidley police log	
ID	Reason
T. Prince	Too slow for traffic
C. Riley	Removal of catalytic converter
O. Wood	Underage driving

Figure 1. Examples of linkage attacks across multiple data sets

2.3 Data Mining In Time Series Data

Power system operators rely on a combination of current and past states (i.e., measurements, topology snapshots, etc.) to obtain the desired levels of observability and control needed to ensure a continuous system operation. As such, the collection and processing of time-series data represents a fundamental enabler of grid applications, with a wide variety of end-applications that range from financial settlements, future demand forecasting to real-time anomalous behavior detectors. Within the energy sector, time series data are often produced or associated with a physical asset or individual, and thus can often be tied to a physical location (latitude/longitude, a building, an operational zone, or another suitable identifier). This allows engineers to assemble and use time-series data in complex spatiotemporal datasets that can be used to reveal and study complex relationships.

Surveys such as those presented by (Atluri, Karpatne, & Kumar, 2018) highlight some of the typical features found on spatiotemporal records, and how these features can be analyzed by using data mining techniques (and hence to extract and analyze private data). According to the authors, spatiotemporal datasets typically contain enough information to reveal: a) **Event data**, which allows identification of a specific activity within a given time or location; b) **Trajectory data,** which can be used to track how an activity evolves during a given time span or how an asset moves; c**) Point-reference data**, a finite collection of data points that can capture a general spatiotemporal behavior (e.g., how irradiance typically behaves during the day) and d) **Raster data**, which relies on the locational attributes to reveal group behaviors or localized trends. Although the above examples assume the inclusion of a locational attribute, non-locational time series can be merged with other time series records to create multi-source, high-dimensional spaces that enable analysts to infer similar types of information.

Due to the vast amount of spatiotemporal data that can be collected, multiple data mining techniques have been proposed in literature to extract hidden information. According to (Atluri, Karpatne, & Kumar, 2018) mining techniques can generally be grouped into:

Clustering: Clustering attempts to group distinct instances within a dataset that share a similar feature or behavior. This technique may be useful in identifying a group of distinct instances that: a) Are located in the same vicinity either on a long-term (fixed assets) or a short-term (such a group of EVs); b) Exhibit a similar activity within a short time-window (e.g., maximum PV output due to a solar-peak); c) Exhibit similar trajectories (e.g., EVs that charge in the same manner, or PVs that have a similar production curve; or d) Exhibit similar time-dependent behaviors (e.g., industries that have similar consumption profiles).

Predictive learning: In predictive learning, a function attempts to learn a mapping between two or more independent datasets. This technique may be useful in: A) Predicting how a time series will behave based on past observations or, B) Predicting the behavior of a feature by observing how other co-related feature vectors behave (e.g., using cloud coverage records to infer solar production).

Change detection: Change detectors seek to identify the moment in time in which a system experiences a permanent change. These techniques may help to: A) Identify abnormal segments or periods in a time series (e.g., to detect curtailment on DERs), or B) Recognize contextual changes, where an individual changes behavior with respect to a community (e.g., when a solar panel within a facility is re-oriented or replaced).

Frequent pattern mining: In the context of spatiotemporal records, a pattern is a series of datapoints that repeatedly appear across multiple dataset instances. Data mining tools could be used to: A) Locate co-occurrence patterns, which are distinct features that share a common trigger or ancestor; B) Identify sequential patterns, in which an initial pattern triggers a secondary pattern; C) Recognize sequential patterns in trajectories, in which different instances follow a similar trajectory (e.g., finding a sequence of EV charging points that are visited during a long weekend); D) Identify motifs within time series, motifs are repeated sequences present in long-term datasets (e.g., to analyze how wind patterns influence in power fluctuations); E) Identify hidden network connections, this can reveal the underlaying communities that individual has joined based on a combination of locality, personal preferences or social facts.

Anomaly detection: Outlier detection algorithms may help to isolate rare or anomalous behaviors. This technique may be useful in a) Identifying data points that break away from previous behavior (e.g., a low degree of autocorrelation), b) Recognizing trajectory anomalies, which can be indicative of a sudden change in behavior (e.g., equipment failure), or c) Identifying changes in raster data behavior, which can occur when multiple, co-located instances suddenly change behavior (e.g., due to disruptive weather events).

Relationship mining: Similarity algorithms may be used to measure the amount of correlation between different dataset instances. These types of algorithms could be used to: A) Identify positive/negatively correlated datasets; B) Measure correlation across time, a technique that can be complemented by inserting a lag function to analyze delayed dependencies.

Clearly, data mining can be used to optimize and enhance many of the day-to-day functions carried by grid operators but could also represent a significant risk to privacy if features can be mapped back to an individual. Ensuring time series data remains useful for operational and decision-making applications while providing sufficient privacy guarantees will require engineers to develop creative solutions that balance both needs.

2.4 Privacy Attacks On Time Series Data

As outlined in the previous section, multiple data mining techniques can be applied to a time-series dataset that could be used to reveal an individual's behavior. Preventing privacy leaks will require engineers to identify, evaluate, and adapt a series of anonymization techniques to thwart potential attacks. These technical evaluations must be carefully designed to avoid faulty implementations that lead to a false sense of security. For example, a typical misconception is to assume that re-sampling, encrypting, or otherwise obfuscating data will protect the underlying data from being analyzed by data mining techniques. An example of such exploitation in time series data is given by (Wright, Ballard, Coull, Monrose, & Masson, 2008), where the researchers analyzed side-channel information to gain additional insight into encrypted Voice-Over-IP (VOIP) services, which could in perfect conditions lead to phrase identification.

A side-channel leakage can occur when a cryptographic algorithm is implemented or integrated in a deficient manner resulting in observable traits that break the perfect secrecy assumptions. For example, in the VOIP protocol, speech waveforms are compressed and encrypted before being sent to the other party. However, for network efficiency reasons, low-amplitude signals are treated as noise and thus get filtered out, resulting in a zero-length package that does not need to be transmitted across the network (see Figure 2). Based on this information, an eavesdropper may decide to monitor the inter-packet timing delay and use it as an indirect mechanism to detect pauses during a conversation (See Figure 3).

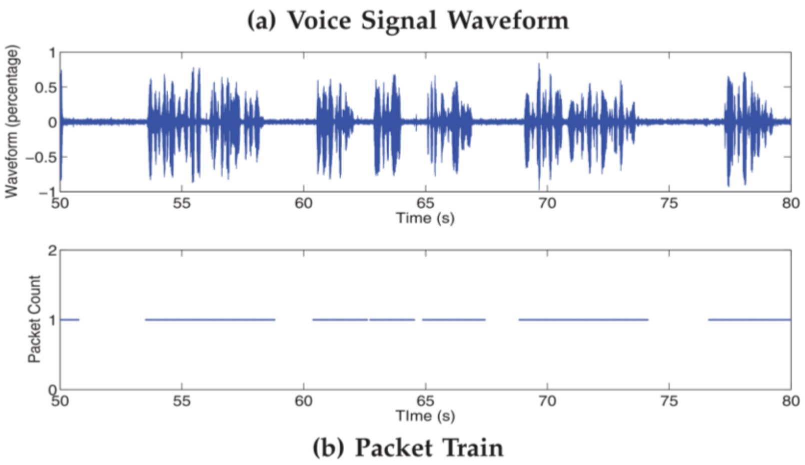
(a) Voice Signal Waveform

(b) Packet Train

Figure 2. Encrypted or packetized waveform data can be used to reveal a presence or lack of data. Graph taken from (Zhu, Lu, & Vikram, 2012)

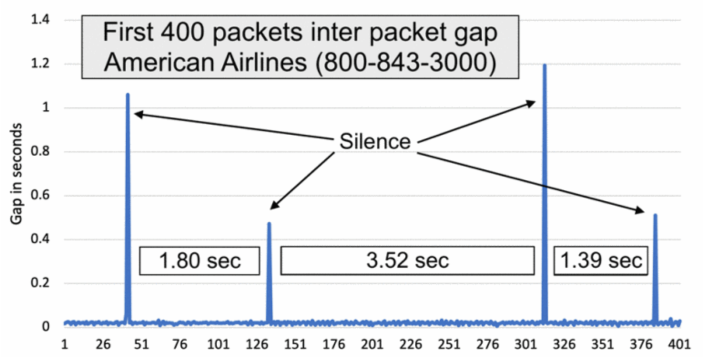

Figure 3. Delays and gaps in a data stream may reveal a change in a stream's contents, in this case silence can be inferred by measuring inter-packet gaps in a packet stream (Laurens, Christianto, Caulkins, & Zou)

Although being able to detect silence may appear irrelevant, the gained knowledge could be used to infer unique behaviors about the individual making the call, such as its cadence, and rhythm. The analysis of side-channel information to gain additional information about the underlaying process has been studied and applied to a wide variety of use cases. For example, in (Tramer, Boneh, & Paterson, 2020), the authors analyzed the response-reply latency in a privacy-oriented financial blockchain to classify if transactions were carried between new or already known parties. Another noteworthy example that illustrates the need for securing datasets that follow a time-based patterns is the de-anonymization attack described in (Narayanan & Shmatikov, 2008), where anonymized records from a video streaming provider were time-correlated with the data provided by a movie-review dataset to reveal the identity of the subscriber.

2.5 Potential Privacy Attacks In Energy Infrastructure, A Review

Within the last decade, multiple researchers have identified and documented a variety of potential weaknesses that could lead to privacy incidents within the energy sector. To illustrate some of the published research Table 1 presents a series of studies that discuss potential vulnerabilities. The presented studies cover a wide range of applications, but in general, they tend to have a high degree of customer interaction, which often results in the handling of sensitive data records. As can be seen in Table 1, many of the attacks rely on merging or combining a mixture of anonymized data records with other publicly available datasets to infer additional attributes about the data producers such as general location, consumption patterns, or mobility patterns. Another significant number of studies consider attempts to remove privacy as a secondary effect that may occur if the primary cybersecurity mechanisms are compromised (e.g., by launching a Man-in-the-Middle attack or gaining network monitoring capabilities). However, as grid operations continue to evolve into a decentralized architecture, it's possible that other risks and targets may appear. For example, grid operators may need to interact with new providers to exchange a variety of grid services, which in turn could translate into a mutual need for information exchange. As such, future data exchanges may contain a variety of field-collected sensor measurements, grid models, or DER states and commands potentially expanding the privacy risk to organizations. Such risks will need to be identified and handled to minimize potential operational or financial impacts. To satisfy this need, a short review of novel techniques that can mitigate privacy risks will be discussed in the next section.

Table 1. An overview of published works related to privacy threats in the electricity sector.

Work	Domain and principle of operation	Overview
(Chen D. S., 2016)	Solar Energy, Data correlation attacks	The researchers presented *SunSpot*, a system for localizing solar-powered homes from anonymized solar production databases. The home's latitude was inferred by correlating day length information (sunset-sunrise) with the recorded production times. The longitude was deduced by matching the solar noon time with peak production timestamps. Publicly accessible satellite images, provided by Google Maps were used to identify homes with solar modules (regions were estimated according to the computed latitude/longitude). Mechanical Turk was used to verify that the address was solar powered. The author's proposed approach demonstrated that location information could be inferred through distinct solar signatures, even when the energy data is anonymous. The authors speculated that address and energy output could be used by malicious actors for criminal activity.
(Badr M. M., 2022)	Edge Energy Management, Inference attacks	The researchers identified that federated learning (FL) model parameters sent to the utility server may be used to extract a resident's private information. The information can be used to infer when homes are vacant, or tenants are sleeping. An adversary can perform a model inversion attack by training a custom ML model on the output of the targeted home's model, essentially removing the protections provided by FL and enabling them to infer the input data. Their model can also enable an adversary to perform a membership inference attack, by determining if a given input was part of the original model's training dataset. Authors conclude that deep learning methods represent an emerging risk to customer privacy due to the vast amounts of data available.
(Razavi, 2019)	Smart meter, Customer profiling attacks	The researchers used machine learning methods operating over smart meter data records to determine the occupancy status of a home. The machine learning model was trained on electricity usage data that was collected over an 18-month period from over 5,000 homes. The usage data was taken from the homes at 30-minute intervals. The researchers found that a weeks-worth of household data was enough to predict a (single) timeslot where the resident was away from their home with 79% accuracy during the following week.
(Fan, 2017)	Smart meter, Customer profiling attacks	The researchers demonstrated a new attack on smart meters that used reactive power data to determine when appliances were turned on or off. They performed experiments on the load profiles collected from three households and found that their attack could accurately detect when an appliance was turned on or off. Appliance-level information can be used to infer additional information about the residents, such as when they are asleep or out of the home.
(Chen D. a., 2017)	Smart meter, data correlation attacks	The researchers developed a tool called *Weatherman. Weatherman* ingests anonymized energy usage and combines them with wind, and solar data to reveal the energy meter location. The authors demonstrated that each location on Earth can be associated with a weather signature that is unique to that area. They tested *Weatherman* against an earlier tool they developed, *SunSpot* (Chen D. S., 2016), and found that *Weatherman* produced more accurate results while relying on fewer datasets and less granular information.
(Hoh, 2006)	Traffic Information	The researchers performed a case study on intelligent transportation systems to determine their vulnerability to inference attacks by

	Systems, data correlation attacks	analyzing the anonymized vehicle position information sent to traffic-monitoring services. By leveraging this information and applying clustering techniques they demonstrated it was trivial to determine the most likely garaging location (e.g., the overnight parking address) for a vehicle in an average suburban scenario. An adversary could further combine location data along with a reverse address database to reveal the driver’s identity.
(Brighente, 2023)	Electric Vehicles (EV), Private/sensitive data exfiltration during a cyber attack	The researchers investigated the methods an adversary could use to obtain private information by eavesdropping on the interactions between the internal EV subsystems and between the EV and the surrounding environment. The EV controller and infotainment system can be used to extract preferences along with driver’s credentials during their energy exchanges. In a situation where wired charging occurs, the interaction between the EV and the charging apparatus can reveal unique EV identifiers that enable an adversary to track the EV. Man-in-the-middle and eavesdropping attacks done on wireless charging stations could allow an adversary to obtain data about billing information and vehicle ID.
(Apthorpe, 2017)	Internet of Things, Private/sensitive data exfiltration during a cyber attack	The researchers demonstrated that a network observer could use network traffic from smart homes to determine home activities. This was possible even when the traffic was encrypted. The proposed attack locates and identifies smart home devices by using the domain name system (DNS) but the author notes that device fingerprinting may also be used. Health conditions could be revealed by analyzing health oriented IoT devices (e.g., a blood sugar monitor). The devices can also reveal the private lifestyles of the residents in the home.
(Huang, 2017)	Sensor networks, Private/sensitive data exfiltration during a cyber attack	The researchers identified security issues that could lead to privacy attacks within the phantom routing process used in Energy Harvesting Wireless Sensor Networks (EHWSNs). The phantom routing scheme relies on a fake source node that acts as a “phantom” node, obscuring the location of the energy harvesting node (i.e., the original source). Authors identified various de-anonymization attacks based on the intrinsic geographic dependency between the source and phantom node’s (which need to be geographically close). Since the node routing path is usually determined by the shortest routing algorithm, an attacker can use this information to follow the path from the phantom node to the destination of the network traffic and backtrack to find the source node location.
(Han, 2016)	Vehicle to grid (V2G), Private/sensitive data exfiltration during a cyber attack	The researchers identified how electric vehicles (EVs) in a vehicle to grid (V2G) network could potentially have private attributes exposed during the charging process. Vulnerabilities in the charging protocols between EVs and charging stations can be exploited by an adversary through attacks such as man-in-the-middle (MITM). Through the MITM attack, an adversary can obtain information sensitive information such as customer name, vehicle ID, and charging location and schedule. A malicious actor can use the location data to perform a malicious action like hijacking the EV.
(Fraiji, 2018)	Electric Vehicles (EV), Private/sensitive data exfiltration during a cyber attack	The researchers identified security issues and attacks that can be performed on Internet of Electric Vehicles (IoEV) deployments. An unauthorized node in the network can listen to vehicle-to-vehicle communications to obtain sensitive information, such as the vehicle’s ID and current location. Malicious nodes could implement continuous monitoring on the vehicle-to-infrastructure across communications

		links to track their location across time. In particular, technologies such as 5G could be prone to this type of eavesdropping attack.
(Kimani, 2019)	Internet of Things (IoT), Private/sensitive data exfiltration during a cyber attack	The researchers identified potential security issues with using IoT devices in a smart grid that could lead to personal data leaks. The authors identified that an increasing demand for IoT products may result in weak security controls that are hard to patch once the product is in the consumer's hands. Collected records could enable an attacker to infer information about electricity consumption trends that can be used to determine if the location is unoccupied. The attacks can also be used to reveal financial information like credit card numbers.
(Pal, 2018)	Microgrids, Private/sensitive data exfiltration during a cyber attack	The researchers described security issues with smart micro-grids (SMGs) that make them prone to privacy attacks. An attacker can eavesdrop on traffic in the network to obtain information about power consumption at various time periods. From this data, an attacker can infer information about the behaviors of the resident and whether the home is currently occupied or not.
(Braun, 2018)	Smart City, Private/sensitive data exfiltration during a cyber attack	The researchers identified a lack of security and privacy defense mechanisms to protect against insider threats within a smart city environment. This may create an opening for an adversary to obtain private data. Exfiltrated data may reveal health conditions, identity, and transportation records (location history). The data may also reveal an individual's lifestyle preferences when intelligent surveillance systems are targeted.
(Shuaib, 2015)	Smart meters, Private/sensitive data exfiltration during a cyber attack	The researchers demonstrated a man-in-the-middle attack in between a smart meter and an upstream server. The attacker's host system manipulated the Address Resolution Protocol (ARP) cache of the meter and server to forcefully route network traffic through their system. The attack showed that transmitted records may be used to obtain private information such as the account holder's name, address, and payment methods.

2.6 Mechanisms That May Help Improve Privacy, A Review

As introduced by Subsection 2.5, multiple researchers have identified potential privacy risks, that if left unattended, could impact consumers and organizations. This has led to the proposal and development of mechanisms that seek to improve upon existing practices with varying levels of success. Some of the published research related to the power industry has been captured in Table 2, from which it can be observed how multiple algorithms have been studied, and evaluated under a variety of uses cases. The algorithms, for the most part, have been adapted from or assembled into more complex processes with the goal to address a specific privacy deficit or potential vulnerability within the use case. Although the specific implementation details vary widely across the literature, some general trends can be observed:

- Protection is generally provided by a mixture of additive noise or encryption mechanisms that seek to distort or hide the information being stored or transmitted across the system. Such approaches may help to add a layer of anonymization to the records, but they may not be able to fully protect an individual if too many features are included in the data set from the beginning.

- For systems based in cryptographic methods, it appears that most researchers are relying on cryptographic primitives whose functional mechanisms are well understood and are thus considered safe. However, adopting a secure cryptographic primitive does not imply a secure solution, and therefore developed applications should be thoroughly reviewed to ensure the algorithms are being used as originally intended, and that no information leakage is occurring (e.g., to prevent side channel attacks).

- Although most of the implementations have been evaluated to ensure functional requirements are met (e.g., processing time, and throughput), they may not be considering the long-term implications of maintaining such systems. This may be range from cumbersome (e.g., providing key management) to the unintentional leakage of private data[1].

In addition to the aforementioned trends, it is also apparent that certain algorithms or methods appear to have fostered a higher degree of interest from the research community. This increased interest may help accelerate the field's overall maturity by allowing researchers to down select approaches based on their technical merits. Based on the research presented in Table 2, it appears that most solutions tend to be driven by the following methods (or variations thereof):

Addition of noise: Adding a controlled amount of noise to individual instances within a dataset could prevent an attacker from learning specific traits about an individual, while enabling statistical analysis at the dataset level (Mivule, 2013). Multiple noise-based methods have been proposed, ranging from the use of white noise that increase the privacy of static records to advanced noise-shaping methods that can guarantee that the presence (or absence) of an individual to a group cannot be inferred. Notable methods include Differential Privacy which will be reviewed in Section 3.

Federated learning: Federated learning is a machine learning method where a central server trains a global model using the parameters provided by locally trained user models (Zhang C. Y., 2021). The typical process works as follows: 1) The server first sends an initial model to users; 2) Users train the model using their local data; 3) The user returns their model parameters to the server; 4) The server aggregates the received parameters to create a global model; and 5) The server sends an updated global model to all users, repeating the process until the global model converges. Additional techniques, such as Homomorphic encryption and Differential Privacy may be applied by edge-located users to their model's parameters, further enhancing their privacy guarantees.

Functional Encryption: Functional encryption is an encryption system that enables a user to encrypt a plain text message (denoted as m) with a user-generated primary key, while also generating a secondary key (Yin, 2021). The second key can be then used by function f to compute $f(m)$ without needing to reveal information about the original m. The purpose of functional encryption is to preserve the privacy of the data contained in m while allowing third parties to verify a non-sensitive attribute (e.g., by enabling a third party to obtain the total money spent instead of individual transactions).

Homomorphic Encryption: Homomorphic encryption allows users to perform a limited set of mathematical operations over encrypted datasets without requiring an intermediary decryption

[1] Refer to Section 4, which discusses theorems that set limits on the amount of noise-modified data that can be shared before it loses its privacy characteristics.

process (Strepparava, 2022). Unlike functional encryption, multiple aggregation patterns (via mathematical functions) are supported by Homomorphic encryption. The result of the computations can be observed by decrypting the new cypher text. The goal of the homomorphic encryption method is to preserve the anonymity of a user's data while still allowing aggregate-level operations to be performed on the data.

Zero-Knowledge Proof (ZKP): A ZKP is a method used between two entities to show that one entity (the prover) knows another entity's secret (the verifier) without revealing the secret itself (Sun, 2021). The method consists of three phases: *witness, challenge, response*. In the *witness* phase, the prover calculates a proof that is given to the verifier; In the *challenge* phase, the verifier generates a series of questions that are intended to assert the prover's true behavior. In the *response* phase, the prover sends its answers to the verifier and a decision is made. ZKPs may also operate in non-interactive manner, where the verifier does not need to ask the prover any questions (by instead relying on a key that contains enough information about the proof itself).

***K*-Anonymity:** A privacy model that guarantees a participant's record is indistinguishable from at least $k - 1$ other participants in the dataset (di Vimercati, 2023). This is accomplished by sanitizing the dataset through generalization and suppression of the data. Larger *k* values provide more privacy to individuals in the dataset at the cost of providing less detailed information. Exact and heuristic algorithms may be used to reduce the amount of data lost during these processes, which may result in the removal of a record if an individual's behavior cannot be hidden by the group. *K*-anonymity is susceptible to homogeneity and external knowledge attacks. A homogeneity attack may occur when an individual can be mapped to a subgroup that is smaller than *K*, this can occur when orthogonal fields are compounded to reduce the search space (e.g., by combining age, gender, and zip code). Methods such as l-diversity and T-closeness seek to mitigate some of these risks.

Table 2. An overview of proposed privacy mechanisms found in literature targeting the grid space.

Citation	Domain and principle of operation	Overview
(Tran, 2022)	Smart meter, added noise to enhance privacy	The researchers demonstrated a privacy system whose goal was to anonymize electricity consumption data collected from smart meters. The proposed mechanism seeks to prevent attackers from inferring the resident's habits and home appliances used. The system used two separate algorithms to generate noise at the smart meter and at the Distribution System Operator (DSO) endpoint. The noise is generated using a private noise distribution protocol called nn-PND.
(Zeng, 2017)	Smart meter, decentralized key generation and management	The researchers proposed a lightweight scheme for collecting meter data in smart grids. The scheme provides a decentralized way for a control center (CC) and users to generate and manage their own keys. A pseudorandom number generator and the Diffie-Hellman key exchange method are used with attributes like CC ID, user ID, timestamps, and other keys to generate new keys. In this scheme, the CC first sends a request to the user to collect their metering report. The user encrypts the report with a

		session key based on the report's timestamp and sends it to the CC for decryption.
(Zhang Y. J., 2018)	Electric Vehicle, Smart Grid, Paillier Cryptosystem	The efficient privacy-preserving communication and power injection (ePPCP) architecture is proposed by the researchers with the goal of keeping an EVs' individual power injections hidden from the utility company. This is accomplished using two secret keys and a hash-then-homomorphic technique. Then the bid is encrypted using a Paillier Cryptosystem based method. The architecture was tested against an adversary model and was found to deter man-in-the-middle and replay attacks.
(Xu S. X., 2021)	Electric Vehicles, Blockchain	The researchers demonstrated how the proposed system could allow the user of an EV to remain anonymous when scheduling, charging, and paying for services provided by the Electric-Vehicle-charging Service Provider (EVSP). The system uses distributed Public Key Infrastructure and blockchain to maintain user and EVSP records. When the user wants to access EVSP infrastructure, they begin a process that involves a zero-knowledge proof, ring signature, and three tokens for scheduling, charging, and payment.
(Kumar, 2020)	Electric Vehicle	The researchers detail an EV charging system that provides increased protection against the leakage of information pertaining to EVs and their location. The system makes use of lattice-based *signcryption* (an integrated signature and encryption process) to deter post-quantum computing attacks. The presented system was shown to obscure the identity of a user by assigning a pseudoidentity to the EV which can later be used at the charging station.
(Mahmood, 2018)	Smart grid, Elliptic curve cryptography, lightweight cryptography	The researchers describe how a lightweight authentication framework used in smart grid communication can deter all known security attacks when based on elliptic curve cryptography (ECC). The authentication framework used also provides privacy to individuals. Proofs are provided for attacks such as man-in-the-middle and impersonation attack. A proof of perfect forward secrecy is also provided.
(Strepparava, 2022)	Smart grid, energy community, homomorphic encryption	The researchers demonstrated how a protocol based on homomorphic encryption could be used to enable a home's production and consumption data to remain anonymous. The protocol was hardware-tested on an energy community, Lugaggia Innovation Community (LIC). The local energy market for LIC was built on a Cosmos based blockchain architecture. To further improve the security of the data, the researchers suggest the use of zero-knowledge proofs to remove the need for an administrator to store keys and thus have knowledge of private information.
(Poh, 2019)	Smart home	The researchers propose *PrivHome*, a framework intended to preserve the resident's privacy by encrypting their smart home data. It uses two protocols: Authenticated Key-Establishment (AKE) and Searchable Encryption (SSE). A secret key is used in the AKE protocol to authenticate all entities involved in the communications and to establish a session key. The SSE protocol uses the session key to generate a secondary key that protects and enables a

		searchable index of the devices on the network. The index is used locate devices through encrypted queries.
(Rabieh, 2016)	Traffic Management, Vehicular Ad Hoc Networks (VANETs)	The researchers demonstrated a scheme for route reporting that prevented an attacker from linking a vehicle's pseudonym in a Vehicular Ad Hoc Network (VANET) with their reported future routes to infer the identity of the driver. The scheme proposed had a variant for VANETs with and without infrastructure. In the infrastructure variant, homomorphic encryption is used. In the non-infrastructure variant, Elliptic curve cryptography (ECC) point addition and homomorphic encryption are used. The non-infrastructure scheme was shown to be capable of thwarting collusion attacks.
(Li, 2020)	Traffic Management, Blockchain, Zero-knowledge range proof	The authors presented a system built on the permissioned blockchain Hyperledger Fabric framework designed to preserve the integrity and privacy of the data collected from drivers (e.g., vehicle ID and location) in a traffic management system. The vehicles and road objects (e.g., toll stations) form a vehicular network in the traffic management system. Each county or state may have their own traffic management system for their area. Zero-knowledge range proofs are used to encrypt the vehicles messages to a gateway that switches them between adjacent traffic management systems. The gateway validation process was built on a cryptographic library provided by Hyperledger Ursa. The researchers showed how the system could deter vehicular data, gateway spoofing, and eavesdropping attacks.
(Lei, 2022)	Renewable Energy, Permissioned Blockchain, Hyperledger Fabric, Homomorphic encryption	Researchers presented an energy trading platform that was designed to reduce the cost and increase efficiency of the trades while providing an increased level of user privacy. The platform was built on Hyperledger Fabric, relying on a channel isolation mechanism to segment users. Only users in the same channel can share data with each other. The Paillier algorithm was used to generate private and public keys for each user in the network. The keys were used to encrypt/decrypt their private data (e.g., account balance) to and from Fabric's ledger. Internet of Things (IoT) technology was used to deploy automated trading processes.
(Badr M. M., 2023)	Federated Learning, Energy Forecasting, Smart Grid	The researchers developed a federated learning-based energy prediction system aimed at achieving high accuracy while preserving the customer's private information. Inner-product functional encryption (IPFE) was used on the parameters of the customer's model so that the data could be sent to the utility provider anonymously. The scheme used by the researchers enabled the utility provider to use the encrypted parameters to build a global model.
(Yin, 2021)	Federated Learning, Multiparty data sharing, Functional Encryption	A federated learning method developed to protect an individual's privacy from being inferred from the training model parameters (e.g., model weights) was presented by the authors. The method used function hiding multi-input function encryption to obscure the model parameters sent to the server. Bayesian Differential Privacy is used to further preserve the privacy of the model parameters.

(Son, 2020)	Smart grid, smart meter, blockchain, functional encryption	The researchers demonstrated an energy trading system made up of smart meters, a distribution system operator (DSO) server, and a private Ethereum blockchain with the purpose of encrypting bids to preserve anonymity of the user's data and transactions. The bid is encrypted with function-hiding inner product encryption (FHIPE). A buyer and seller are matched with functional encryption-based smart contracts applied to the bids.
(Kiarie, 2019)	Smart meter, Encryption methods	The researchers compared Spritz vs Rivest Cipher 4 (RC4) as encryption methods for protecting electricity consumption records collected from a single home. Unencrypted data could be used by an adversary to infer personal information such as lifestyle choices and sleep patterns. Performance tests (i.e., delay overheads) were carried to determine if Spritz could replace the RC4 encryption specified by the smart meter standard. After 10 testing rounds, Spritz was on average 1.889 times slower than RC4, but didn't prevent the smart meter from functioning normally. However, Spiritz took 3.21 times longer to brute force when compared to RC4.

3 Evaluation Of Privacy Constructs In The TES Domain

A Transactive Energy System (TES) is defined as a set of *"Economic and control mechanisms that allows the dynamic balance of supply and demand across the entire electricity infrastructure using value as a key operational parameter"* (GridWise Architecture Council, 2018). Such a system could be used to:

- Enable the deployment and operation of autonomous agents that work collaboratively towards achieving a common set of goals.
- Enable individuals to contribute and benefit from a shared pool of resources. The collaborative approach may allow individuals to reach efficiencies, or resilience levels that would be unattainable at the individual level.
- Enable the fast and reliable integration of Distributed Energy Resources (DERs) into everyday operations by abstracting the complexity of managing resources in a 1:1 relationship.

Clearly, a well-designed TES can be used to support and deliver a wide array of benefits to their members. However, this also requires a constant exchange of data among participants to signal their internal states and future intentions. Although the specific data being exchanged could vary widely among the participating entities, typical data flows that could be exchanged are:

- **Resource availability information:** This data may be used to reduce the uncertainty levels associated with a variable resource. For example, an entity may need to announce a temporary reduction on storage capabilities to satisfy a local need.
- **Operational states:** Individual entities may share current as well as planned demand levels with neighboring nodes to maximize efficiency. In a TES environment, this may also include the sharing of operational data (switch status, voltage, power flows) or other variables that are typically associated with Distribution System Operators (DSOs).
- **Control actions**: This may include price signals, and other direct control actions (e.g., switch operations) that denote an intention to transition into a new state.

Although it would be possible to have a centralized communication architecture that can effectively isolate data flows on a need-to-know basis (and thus eliminating many of the privacy risks), such approaches would break from the intended goal of moving towards a more decentralized grid that can adapt, and automatically reconfigure when necessary. Therefore, it becomes necessary to explore solutions that can continue to provide privacy while allowing data to move freely across a service region, enabling individual agents to independently observe and assess their local environment but without compromising the privacy rights of their neighbors.

Based in this principle, this section explores the use of two privacy mechanisms to identify their strengths (and deficiencies) in a decentralized environment. Although the evaluation context is limited by the needs of TES-based deployment, the goal is to enable system designers to determine if the proposed methods are suitable for their application.

3.1 Differential Privacy, A Primer

As discussed in Section 2, multiple anonymization techniques have been proposed to secure grid-related data. A commonly used family of methods is based on the introduction of noise to perturb individual records, with the goal of enabling privacy at the individual level while enabling data analysis at the population level. One of the methods that has received a significant amount of interest within the last decade is Differential Privacy (DP). DP achieves its protection characteristics by adding a controlled amount of noise in such a manner that an attacker cannot determine if an individual is (or is not) a member of a given population. The amount of added noise is an application-defined parameter that must be defined based on the privacy needs versus the desired accuracy.

Differential Privacy is based on the *randomized response* survey model first introduced by (Warner, 1965), and refined by multiple authors across the years (Nayak, 2020). In a randomized response, the individual under study uses a randomizing device (e.g., a fair coin) to choose whether a truthful answer or a *random response*[1] should be used. Since the subject under study operates the randomizing device without communicating its output, the surveyor cannot assert the subject's true answer, however, they can still infer population characteristics by removing the statistical bias introduced by the randomizing device. An example of a randomized response process can be observed in Figure 4, where a subject starts by flipping a coin. If the coin lands on heads, it answers the question truthfully, if the coin lands on tails, the subject must flip the coin again and use the result to answer the question. Such a process offers the subject the ability to deny any recorded answer (and instead assert it was a random response), while still producing meaningful results at the population level (e.g., the effects of the coin flip can be cancelled out).

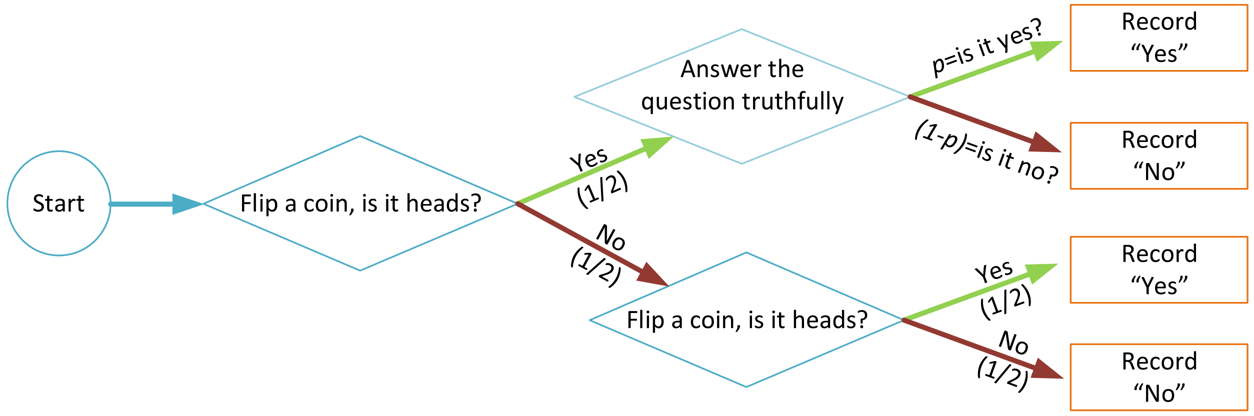

Figure 4. The probability tree in a randomized response survey, assuming a fair coin is used.

Based on the probability tree presented in Figure 4, it can be shown that the expected number of true "Yes" answers can be modeled by $Pr(yes) = \left(\frac{1}{4}\right)(1-p) + \left(\frac{3}{4}\right)p = \left(\frac{1}{4}\right) + \frac{p}{2}$. Therefore, for a significantly large population, P can be estimated as $2\left(\frac{\#yes}{pop.size} - ¼\right)$.

[1] In this context, *random response* is used to describe a response that is not related to the question under study. As such, its value could be determined by a random process (e.g., a coin flip) or be the response to a non-sensitive question, such as *Do you like apples?* whose distribution is known *a priori*.

Differential Privacy requires the presence of a *trusted curator*, who receives raw data from data producers, applies a data transformation algorithm, and delivers privacy-protected results to data consumers. Curators can reply to data consumers requests in an iterative manner or produce a static, artificially generated dataset that captures the behavior of the population under study. Differential Privacy can be more formally defined by (Dwork & Roth, 2014):

$$Pr[M(x) \in S] \leq exp(\varepsilon)\, Pr[M(y) \in S] + \delta \qquad (1)$$

Where:

$\Pr[c]$ is the probability associated with a given condition, for example $Pr(X = x)$ is the probability that the random variable X is equal to the value x, similarly $Pr(W \in Z)$ is the probability of a random process W will output a value in set Z

$M(a)$ represents a randomized algorithm operating over input space A (e.g., $a \in A$) that yields $M(a) = b$ with $Pr(M(a))b$ for each $b \in B$

x and y are similar databases who vary by at most 1 record, e.g., the same subject is either present or not. Such databases are often referred to as neighboring databases.

S represents the output domain of the DP function. This implies that x and y must have similar features (e.g., they need be related) to allow $M(a)$ to compute a valid result within the DP's output domain.

$Exp(x)$ represents the exponential function, i.e., e^x.

ε represents a measure of the desired privacy level and is often called the privacy parameter or privacy budget. Small ε values will yield better privacy (at the cost of less accurate responses)

δ represents a small "failure probability" to handle cases where $exp(\varepsilon)$ does not hold (and thus there exist a possibility of releasing unprotected data)

A formal reasoning on DP principles and mechanisms can be found in (Dwork & Roth, 2014). However, an intuitive explanation of how Eq. (1) operates can be described as follows:

- Assume x and y are two similar databases who at most vary by 1 record.
- Assume that there exists a randomized algorithm M that when applied to dataset y it will yield a result that is very close to $M(x)$, thus making it difficult for an external observer to identify if x or y was fed to M.
- Assume that $exp(\varepsilon)$ can be approximated by $1 + \varepsilon$ for moderately small values, ε its defined by the application's privacy needs.
- Then it follows intuitively that it is possible to express the relationship between $Pr[M(x) \in S]$ and $Pr[M(y) \in S]$ as a ratio, and introduce a restriction to bound it to a desired ε-differentially private level:

$$\frac{Pr[M(x) \in S]}{Pr[M(y) \in S]} \leq (1 + \varepsilon)$$

- Finally, by introducing δ to handle the imperfect nature of a real-world M implementation, we could then define the relationship between both sides as:

$$Pr[M(x) \in S] \leq exp(\varepsilon)\, Pr[M(y) \in S] + \delta$$

It's important to note that DP does not dictate the mechanisms responsible for generating the randomized output of algorithm M, it is up to the system designer to select the best algorithm given its needs. However, some of the most commonly used mechanisms include:

The Laplace mechanism: As its name states, it generates an output value by sampling the Laplace distribution. The Laplace distribution is generally used to model independent outcomes that have exponential behaviors, such as income distribution. This mechanism is often used to protect records that can be counted (e.g., quantities)

The Gaussian mechanism: The mechanisms samples from the gaussian bell, often resulting in the introduction of a non-zero δ-valued parameter. However, it is particularly useful when multi-step queries need to be implemented while controlling the amount of noise being added. For example, the query *"compute the average load for high-demand consumers",* requires the use of a *"greater than"* filter, a "summing" function and a "counting" function to perform:

$$Average = sum(where(x_i > value))/count(where(x_i > value))\ .$$

The Exponential mechanism: This mechanism returns an unperturbed data record from a database that is *likely* to satisfy the original query, but not necessarily the best match (hence protecting privacy). This could be particularly useful to implement queries such as "*Select a DER to be curtailed who has not been curtailed in the last 7 days."*

In addition to the fundamental goal of DP, which is to encourage individuals to participate by providing a mathematical guarantee over their anonymity within a dataset, well-implemented DP-based methods should be able to (Dwork & Roth, 2014):

1. Provide users with a wide protection against arbitrary risks, not only providing protection against re-identification attacks.
2. Prevent the use of linkage attacks, including those that attempt to use a combination of past current, and past datasets to increase their knowledge.
3. Enable the quantification of privacy by defining a fixed bound on the amount of privacy loss. By defining a fixed bound, multiple algorithms can be evaluated and be compared for accuracy or speed.
4. Allow algorithmic composition: Different DP processes can be assembled in series or parallel while still being able to compute the net privacy loss. In a series composition, in which different DP methods are daisy-chained (e.g., when the output of a stage is fed into another step) the privacy loss is the sum of the individual DP mechanisms[1]. In parallel composition, in which multiple DP methods are applied to disjoint sets in a database (e.g., by implementing parallel batch processing) then the privacy loss factor remains constant.
5. Adjust the privacy budget, so it becomes possible to fine tune the allowable "distance" between two neighboring sets. This could be used to account for group behaviors.

[1] Note that this value represents the worst possible case. In practice, advance composition techniques have shown to yield better performance (Mironov, 2017)

6. DP protected records offer *closure under post-processing*: This term is used to indicate that it is infeasible to develop an algorithm that makes a result less differentially private, even when additional auxiliary data is available.

Despite all the potential benefits associated with Differential Privacy, DP is not a perfect solution and may fail to protect individual records when hidden correlations exist between two or more "apparently" independent dataset members. For example, in (Almadhoun, Ayday, & Ulusoy, 2020) the authors exploit genetic markers that are typically inherited among family members to extract and deanonymize their records in a genomic database. In the context of power systems, similar scenarios may appear when two or more apparently "independent" records are derived from the same data collection point or are being double counted. Such attacks are often referred as the "tuple dependency attack" (Liu, Chakraborty, & Mittal, 2016) and may break the guarantees provided by DP.

3.2 The Importance Of Time Series Data In TES Applications

As stated in the introduction, TES require a strong communication backbone that can support the communication needs of the developed solution. According to previously published research (Gourisetti, et al., 2021) technologies such as Distributed Ledger Technology (DLT) could be used to build TES solutions that are highly secure, decentralized, and scalable. A DLT is a network of independent nodes that can achieve distributed consensus among fully autonomous agents by using fault-tolerant protocols. DLTs offer tamper-evident storage capabilities as well as smart contracts, which are logical pieces of code that help automate and validate interactions between parties. The self-sufficient nature of DLT networks, allow independent agents to conduct transactions without the need for an intermediary or explicit trust anchor which can be ideal for enabling next-generation, fully decentralized grid interactions.

Although decentralized technologies such as DLT offer an array of potential benefits to system designers, they often require an increased amount of data exchanges when compared to centralized solutions. In addition, these data exchanges may require the involvement of parties who would otherwise be excluded in a traditional approach. Such involvement can help increase the security of systems (e.g., by eliminating single points of failure, or enabling third party validations), but could also introduce new risks, such as enabling the unauthorized monitoring of data streams. Within the context of DLT systems, some of the aforementioned risks may be addressed by implementing ledger-level segmentation, which can help isolate traffic by effectively creating subsystems where actors can be grouped according to their role, geographical influence, or other relevant attributes. However, the number of channels is finite, limiting its applicability in large systems, such as the electric grid. In addition, segmentation can be hard to maintain (or infeasible) in use-cases where actors are involved in multiple, parallel roles or when role transitions occur often.

Based on these limitations, it may be helpful to think that communications in DLT-based solutions (like many other decentralized solutions) are effectively viewable by all system participants, hence requiring the use of novel techniques that can help guarantee privacy regardless of the implementation-specific features that an individual solution may provide. With this goal in mind, the use of Differential Privacy in the protection of time-series data will be evaluated in this section. Although the need for protecting time-series data may not be immediately apparent to TES solution developers, the risk appears when malicious actors are assumed to be present in a TES deployment. These ill-intended actors may choose to (privately) collect system snapshots and use them for posterior analysis, hence assembling a time-ordered sequence of events that may be used to analyze another actor's behavior. In case of DLT-based systems, the ledger itself could contain enough information[1] to reveal a user's past actions, and thus be used to infer behaviors.

Although the specific nature of the data being exchanged by a TES implementation will dictate the type of information that a malicious node can collect, a market-based TES system may produce enough information to reveal or infer:

- An agent's response to price signals, which could reveal information about its cost sensitivity.
- Production and demand values, which may reveal the types and the characteristics of the installed assets (e.g., capacity, ramping rates) and their operational states (e.g., offline, charging, or discharging)
- Bidding information, or other market artifacts to reveal the economic strategies or goals of the entity operating the asset.
- Scheduling information which may help reveal resident's behaviors (being inside or outside)

Naturally, the amount of information available for analysis will be dependent on the number of parties involved, the data being exchanged, the rate at which is produced, and the ability of a given actor to collect and store such data. However, to illustrate the potential risks associated with obtaining access to high-quality data, access to 15-minute resolution consumption data (as typically available from smart meter infrastructure) was used to evaluate DP-based methods' ability to protect privacy.

3.3 Evaluating Differential Privacy In Time Series Data

Currently, access to smart meter data is restricted to the utility providing service, who occasionally offer data access to the consumer being metered. Although energy metering is a key enabler for revenue collection, it is a relatively trivial task that does not require access to a high-sampling rate. Grid operators often configure relatively high-sampling rates to facilitate integration with other high rate-sampling subsystems thereby helping to enhance system visibility and controllability. However, as an indirect benefit from the increased sampling rate, certain technologies, such as Non-Intrusive Load Monitoring (NILM) tools can be used to enhance visibility into assets located behind the meter.

At their core, NILM-based algorithms perform load identification by executing a signature comparison algorithm that contrasts aggregated load records against previously characterized asset signatures. As such, multiple implementations have been proposed, each of them relying on a wide variety of input features, such as Voltage (V), Current (I), Real Power (P) or Reactive power (Q) to analyze frequency-domain patterns, quantity relations (e.g., P vs Q), or transient signatures. The comparison algorithms may rely on Hidden Markov Models (HMMs), Graph Signal Processing (GSP), and deep learning constructs to perform the actual signature comparison (Faustine, Mvungi, Kaijage, & Kisangiri, 2017). Although each algorithm has unique performance characteristics, in general, the accuracy of NILM methods is dependent on A) The sampling rate, which usually ranges from sub-second to daily snapshots; and B) The asset type, which can range from small appliances, HVACs, EVs, to DERs. Higher accuracies are often achieved when higher sampling rates are available (≤ 15 mins) or when large or long cycling appliances are used (EVs, HVACs) (Teng, Chhachhi, Ge, Graham, & Gunduz, 2022)

[1] Although specific DLT implementations may expose their ledgers publicly (e.g., the open internet), the authors assume the presence of a permissioned DLT, which requires all participants to have an organizational-level identity and credentials before access is granted to the application's ledger.

Clearly, due to their fundamental nature, NILM algorithms are well suited to find patterns in noisy data. Based on this observation and the known limitations of DP, it was hypothesized by the authors that DP may not be the best approach for protecting demand data, even at aggregated levels, especially if we assume a malicious agent is capable of time-aligning individual snapshots. Based on this hypothesis, the following subsections are used to evaluate DP's ability to protect aggregated time-series data.

3.3.1 Experimental Overview

Differential Privacy is known to exhibit weaknesses when correlated data is present. Limitations on DP when applied to time series data have been documented in works such as (Liyue & Li, 2014), (Xu H. W., 2017). However, electrical demand data has additional properties that make it particularly difficult to protect, these are:

High degree of autocorrelation: Generally speaking, load patterns are highly repetitive. Load patterns re-occur on a weekly basis (weekdays vs weekends) and have seasonal tendencies (e.g., ACs operate in summer and heating systems operate in winter).

Correlation with external systems: Weather, which data is widely available at the locality level, is highly correlated with HVAC operation, as well as wind and solar production. In addition, these systems represent a significant part of a consumer's net demand (or production) curves.

Systematic data collection: Due to the operational nature of power systems, data is often collected at fixed intervals from multiple locations. This enables time-aligned and locational grouping of multiple, independent data collection points to assemble richer data sets.

During the rest of the section a technical evaluation of DP as applied to consumer-level records will be discussed. The results indicate that DP offers a limited set of privacy protection guarantees when used to protect long-term, demand data even in aggregated form. The resulting privacy leaks could be used to reveal behaviors of individual actors.

3.3.2 Technical Evaluation

NILM algorithms tend to perform better when high-quality, long-term datasets are available for analysis. However, as outlined earlier, access to high-quality, customer-level, demand data is restricted by most utilities. Therefore, research often relies on datasets that are derived from laboratory environments that may fail to fully capture the day-to-day behaviors present in real-world conditions. Cognizant of these limitations, various research organizations have launched field deployments that collect and aggregate data from volunteers across a service region. Collected records may then be post-processed to ensure quality and be further anonymized to prevent their misuse. To illustrate such a record, a 5-minute demand curve, which has been averaged using a 7-day rolling average is presented in Figure 5. The data was originally downloaded from the utilities' customer portal, and later donated by a member of the team. The unfiltered data was collected during the 2022/2023 Winter season by using a revenue-grade meter (with a stated accuracy of 0.5% or better). The site operates in a demand-only mode, with the largest load being a 14.4 kW electric furnace.

To illustrate the behavior of DP, a Laplace mechanism was applied to the aforementioned dataset to enhance its privacy posture (see Figure 6). The result was computed by using the library available at (Lets Make It LLC, 2020), the privacy budget was set to a unitary value ($\varepsilon =$

1). Note that selecting an appropriate value of ε requires a good understanding of the data and defining an acceptable risk threshold. Such discussions are outside the domain of this report but can be reviewed in (Lee & Clifton, 2011).

As demonstrated by Figure 6, DP has introduced a significant amount of noise that should prevent attacker's from extracting individual daily records from the dataset. In some instances, the added noise resulted in negative values which must be removed during post-processing to avoid misclassifying the consumer as a prosumer. However, upon closer inspection, it becomes evident that time-based patterns continue to exist. For example, a repetitive sudden increase in load, indicative of an "*off-to-on"* transition appears near the 3:00 am mark (see Figure 7).

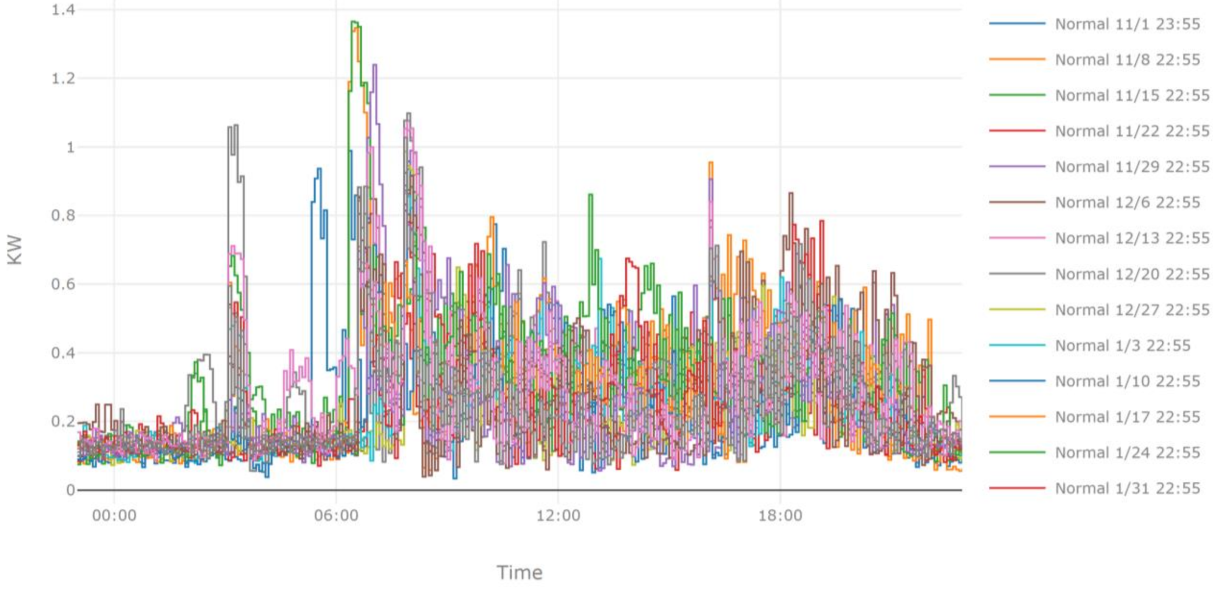

Figure 5. A 7-day rolling average applied to a 5-minute demand curve downloaded from a utility's customer portal.

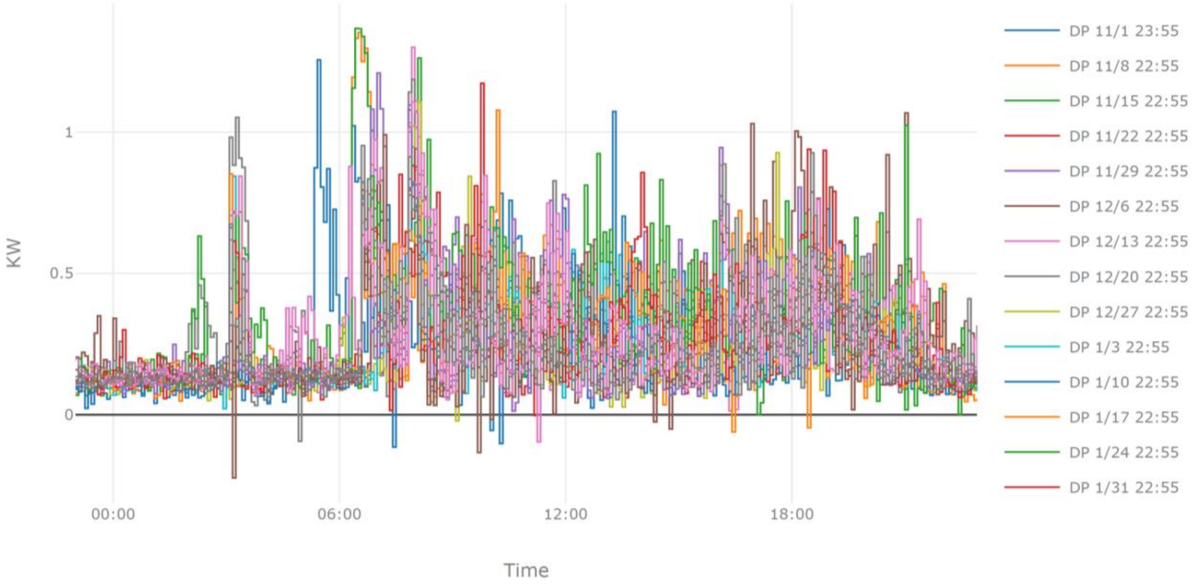

Figure 6. The Laplace DP mechanism as applied to the consumption records shown in Figure 5 ($\varepsilon = 1$).

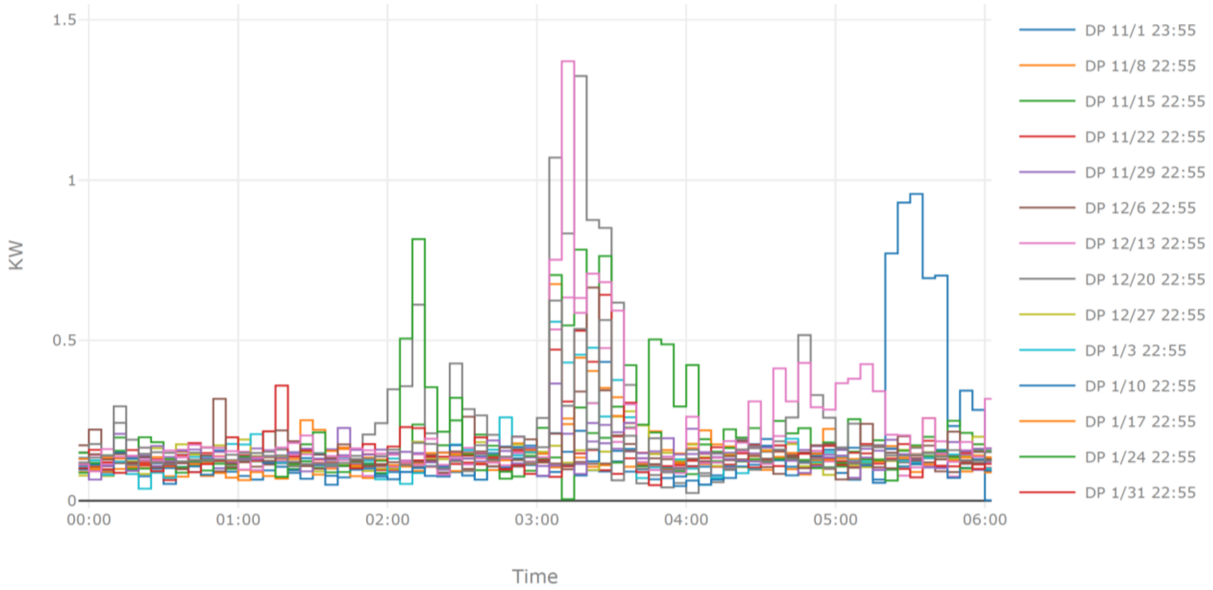

Figure 7. A zoom-in into the plot shown in Figure 6, highlighting the presence of an "event" near the 3:00 am mark.

Such "patterns" could be identifiable by human actors and be complemented with other sources of information to infer individual habits. For example, the subject's "active" period could be inferred to be from 6:30 am to 11:00 pm by looking at Figure 6 (and by assuming higher loads are indicative of an active subject). To better characterize the ability of DP to protect against these inferences, a test was devised to quantitively measure DP's pattern protection performance using a NILM-like approach. The test consisted in training and evaluating the accuracy of a NILM model operating over two distinct time-series, with one of them been privatized and the other being a raw, non-privatized dataset. Such evaluation requires the existence of a "ground truth" signal that describes the state of a given appliance, and a comparator that scores the accuracy of the method in predicting its state.

Based on the test goal, an input data set was procured from the End Use Load Research (EULR) project, a project sponsored by the Northwest Energy Efficiency Alliance. The project aims to support the realization of clean energy goals by enhancing the utilities' ability to characterize the operation of residential-scale heating and cooling systems (Northwest Energy Efficiency Alliance, 2020). To achieve this goal, the EULR project collects equipment-level and net energy consumption data at 1-minute intervals for hundreds of residential homes located within the Pacific Northwest. Although the project charter focuses on collecting demand profiles for different types of heating and cooling equipment, the platform also captures solar production and Electric Vehicle charging profiles when available.

Access to EULR's 1-minute data is restricted to project participants, but aggregated demand values recorded at 15-minute intervals are publicly available. Due to the amount and diversity of measured equipment, as well as the collection period (over 3 years as of 2023), the data set provides an ideal repository for performing an evaluation of NILM algorithms. Furthermore, by having access to asset-level demand, it becomes possible to establish a ground truth that can be used to quantitively assess the effectivity of DP against pattern identification algorithms.

3.3.3 Results

As part of the testing phase, a convolutional neural network (CNN) implementation was chosen for the NILM module. The selected library is based on the work performed by (Zhang, Zhong, Wang, Goddard, & Sutton) and is available at (JackBarber98, 2019). As described earlier, two tests were carried out to determine the effects of using DP on a NILM system. Case A assumes that a malicious agent has access to the unprotected 7-day rolling average demand for a particular customer, while case B, assumes that consumption data has been DP-protected before being compromised. In both cases, it has been assumed that the malicious agent can independently obtain an asset's consumption signature and thus has been left unprotected. The CNN-based algorithms used the 55% of the aggregated demand records for training, 20% for training validation, while the remaining 25% were used for performance evaluation (e.g., 75% of the records were used in training, 25% for blind testing).

The results of test A and test B can be observed in Figure 9 and Figure 10 respectively. Note that the selected NILM algorithm outputs a load factor for a given appliance (circuits labeled as heaters were used for this experiment). However, this functionality was ignored to simplify the evaluation and instead the comparator measures the number of times the device is correctly labeled as being in the "On" state. The "On" state can be configured via a threshold value to avoid false detections.

The variable "threshold" functionality was used to assemble Table 3. As it can be observed the accuracy varies depending on the selected threshold, for DP-protected records, there is a decrease in accuracy as the thresholds increase. The final two rows of Table 3 are used to report the average accuracy, note that the CNN algorithm successfully demonstrated its ability to make high-quality predictions (Accuracy >90%) when unprotected data records are used, this accuracy drops to ~60% when DP-protected records are used.

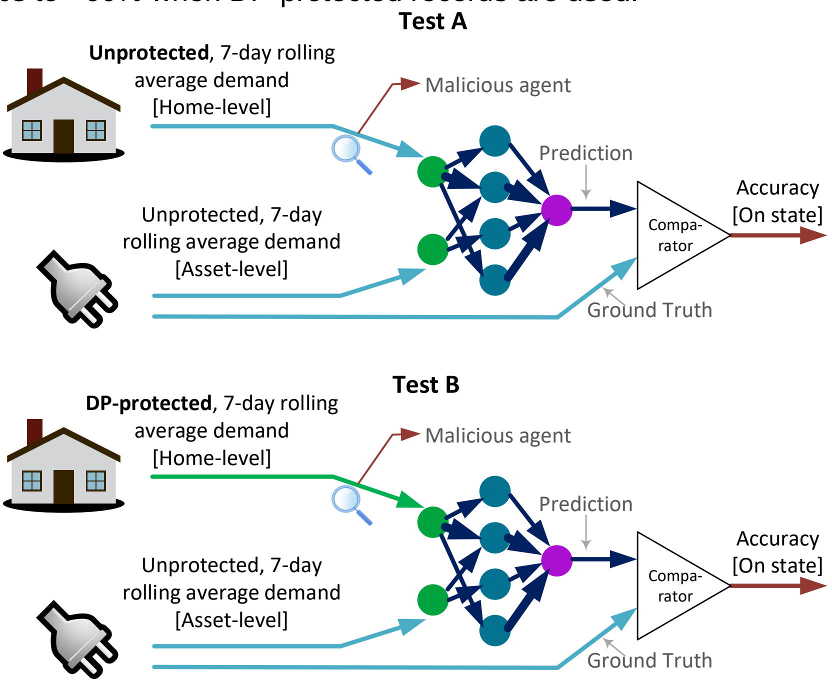
Test A
Unprotected, 7-day rolling average demand [Home-level]
Malicious agent
Prediction
Unprotected, 7-day rolling average demand [Asset-level]
Compa-rator
Accuracy [On state]
Ground Truth
Test B
DP-protected, 7-day rolling average demand [Home-level]
Malicious agent
Prediction
Unprotected, 7-day rolling average demand [Asset-level]
Compa-rator
Accuracy [On state]
Ground Truth

Figure 8. An overview of the testing methodology. The CNN-based NILM algorithm was trained, and subsequently evaluated on A) a Non-protected dataset and B) A DP-protected dataset.

The test results indicate that DP may limit the accuracy of NILM-like processes, hence deterring its use, but may not be able to offer perfect protection (a 60% accuracy remains better than a coin toss). Furthermore, the amount of added noise may severely impact a system's ability to use the data in other operational scenarios (for comparison purposes see the blue-colored lines present in Figure 9 and Figure 10)

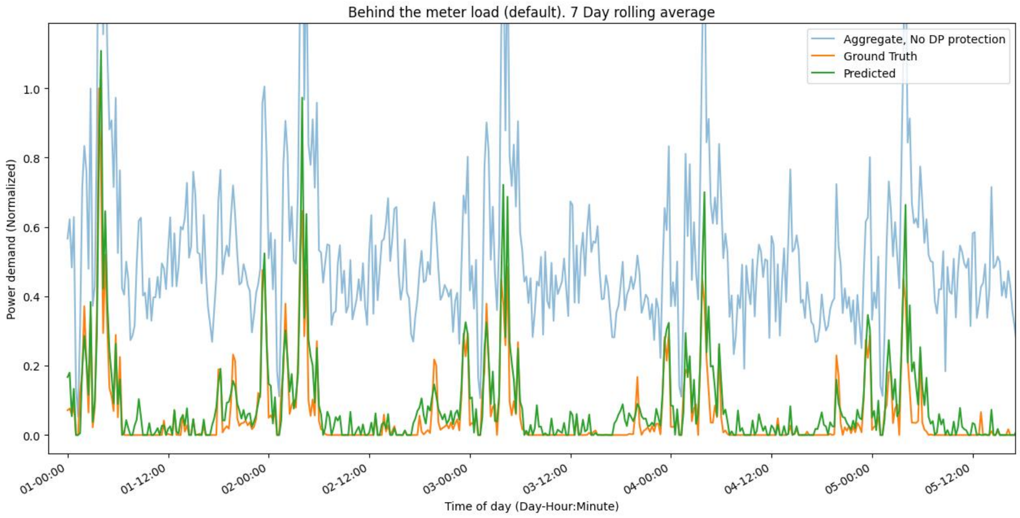

Figure 9. A CNN-based NILM algorithm applied to a 7day rolling window dataset with no DP protection.

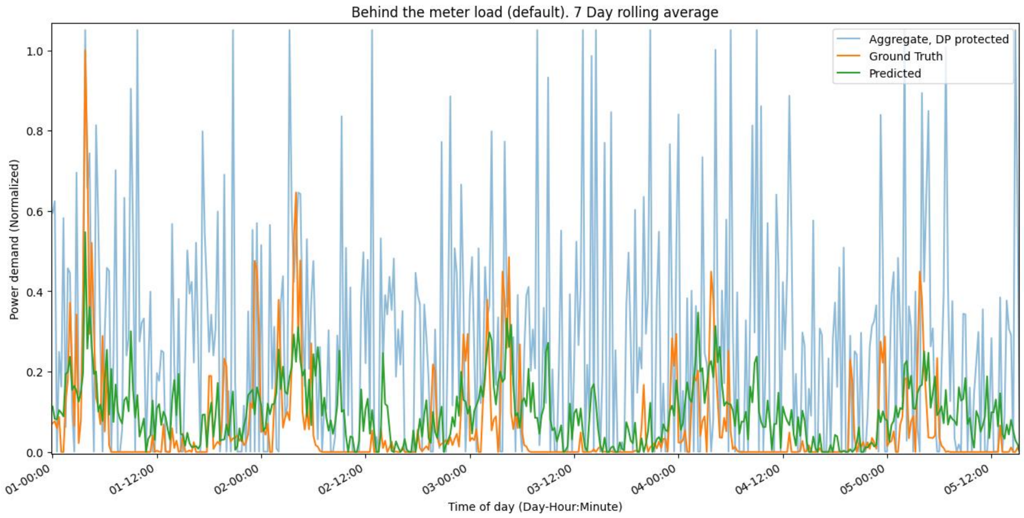

Figure 10. A CNN-based NILM algorithm applied to the same 7-day rolling average dataset with a DP function applied to the aggregate signal.

Table 3. NILM’s ability to predict an appliance’s “Active” state, with a variety of thresholds.

On/Off Threshold	#Ocurrences [Ground Truth]	Accuracy, No-DP protection [%]	Accuracy, with DP protection [%]
0	715	89.65	95.94
0.01	539	91.28	94.43
0.02	446	91.26	93.05
0.03	369	93.22	93.22
0.04	313	93.29	91.37
0.05	283	92.23	87.99
0.06	243	90.95	83.95
0.07	226	90.71	84.07
0.08	211	90.52	81.99
0.09	199	89.95	76.88
0.1	183	89.62	75.96
0.11	172	88.37	74.42
0.12	159	88.05	71.07
0.13	148	88.51	65.54
0.14	144	86.81	61.11
0.15	140	87.86	57.14
0.16	139	88.49	53.24
0.17	133	89.47	54.14
0.18	118	90.68	53.39
0.19	107	90.65	50.47
0.2	103	90.29	46.60
0.21	97	90.72	46.39
0.22	76	98.68	47.37
0.23	62	96.77	46.77
0.24	59	100.00	42.37
0.25	55	100.00	41.82
0.26	50	100.00	40.00
0.27	47	100.00	38.30
0.28	43	100.00	39.53
0.29	38	100.00	39.47
0.3	32	100.00	43.75
Average Accuracy		93.66	60.64
Average Accuracy (Weighted)		91.24	63.07

3.4 Evaluating Differential Privacy In Aggregation Tasks

As discussed in Section 3.1 Differential Privacy provides strong guarantees against a wide variety of attacks when tuple-dependencies are eliminated and appropriate ε values are selected. Such conditions could exist within energy applications if: 1) Sufficiently large population pools are assembled; 2) Records are collected over predefined windows of time based on risk tolerance[1]; 3) Appropriate ε values are chosen based on the amount of noise tolerance versus the desired privacy levels. For example, a days' worth of aggregated demand data from a small-sized group (e.g., customers attached to a pole-mounted transformer) may be safe to share because it's hard to generalize a user's behavior from a single observation. Similarly, feeder level data could be collected and shared over a much longer time-span due to the vast amounts of data required to successfully train disaggregation algorithms. To illustrate this behavior, a subset of the EULR's dataset was mapped to the IEEE 123-Transactitive system by normalizing the EULR's records and using them as the demand curve shapes within an OpenDSS simulation. This process ensures that each load maintains its rating, while inheriting the realistic demand pattern captured within the EULR's database (See Figure 11).

To better understand the implications of applying a DP function over short windows of time, a graphical comparison of applying a non-protected versus a DP-protected summing function over individual consumers attached to Bus 107 (of the IEEE-123 Transactive system) is presented in Figure 12. As observed in earlier cases, the added noise introduces a significant amount of error that may exceed the end user's tolerance expectations, however it makes it extremely hard for an attacker to extract an individual's demand curve without access to external information. Similarly, Figure 13 aggregates feeder-level demand records by using a non-protected and a DP-protected summing function, however, due to the vastly larger aggregation size, the error becomes insignificant.

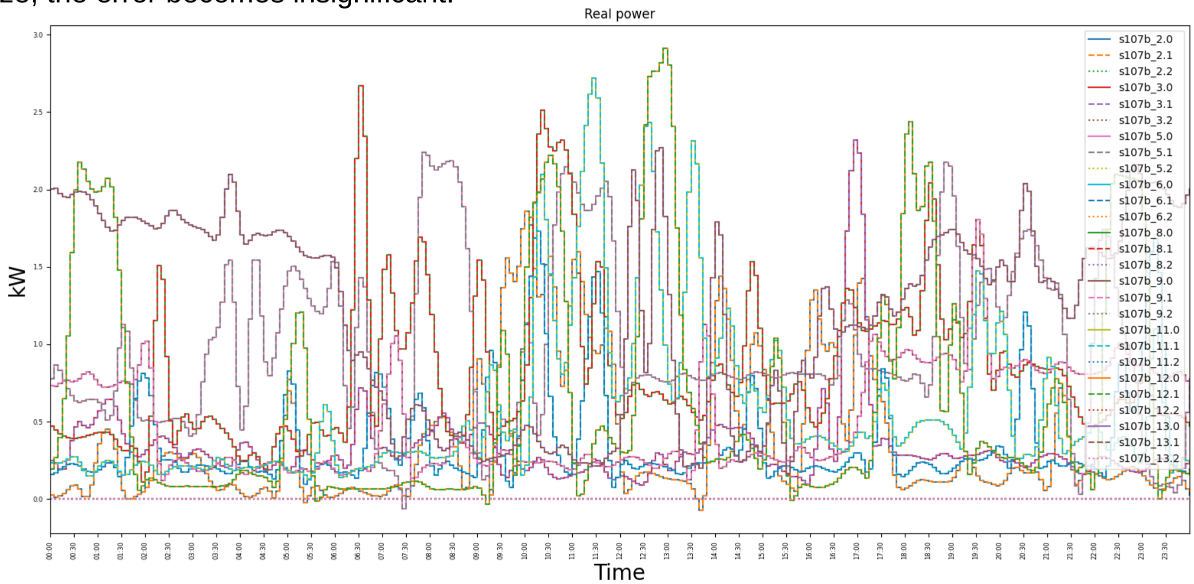

Figure 11. Individual consumption patterns of customers attached to a pole mounted transformer (Bus 107 in the IEEE 123 bus system)

[1] In addition to limiting the window of time over which records are collected, it would be useful for systems designers to identify a dataset's privacy requirements as a function of time. For example, the release of live data may pose a larger risk than releasing old records.

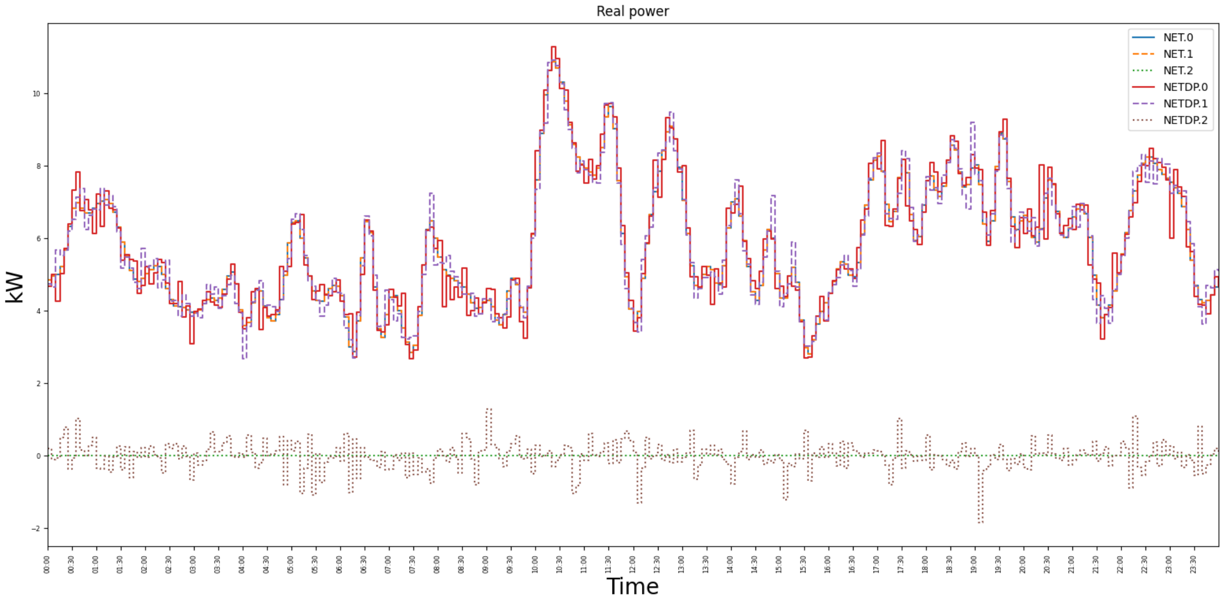

Figure 12. Aggregated demand at the transform-level (from 9 distinct consumers). The amount of introduced error may exceed the application’s expected tolerance.

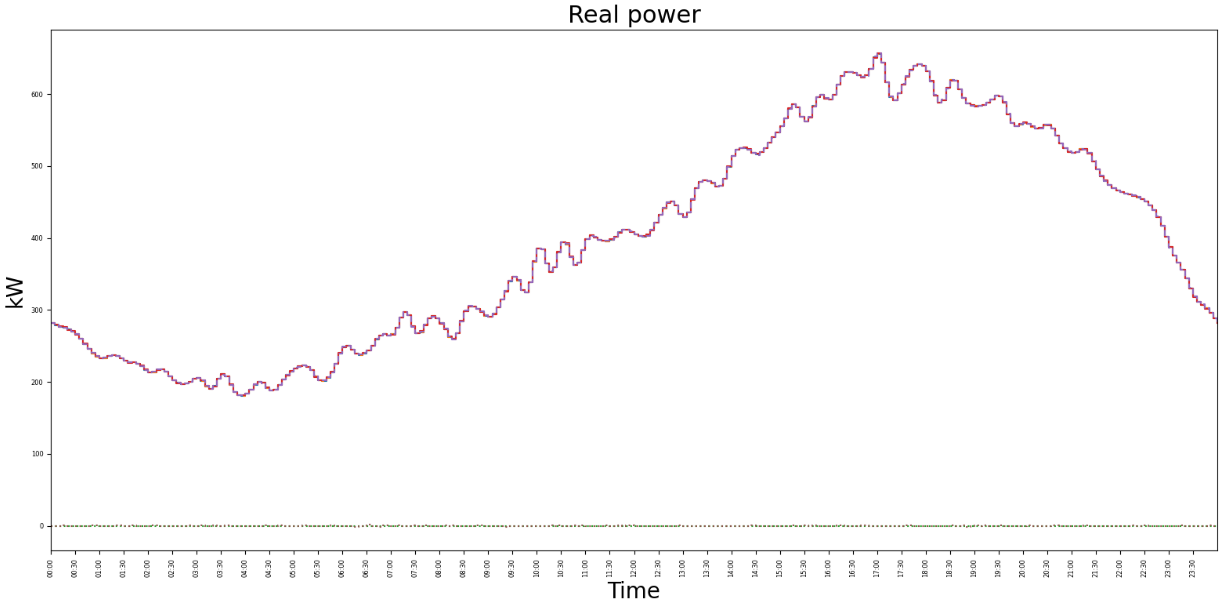

Figure 13. Aggregation at the feeder level, demonstrating the negligible amount of error introduced by DP.

Based on the aforementioned observations, Figure 14 was constructed to illustrate the amount of noise introduced as a function of the population size when DP methods are used. The figure was constructed by aggregating a variable number of demand curves, and then measuring the average amount of relative error present within each aggregation round. The results clearly show that DP is a viable tool when a sufficiently large pools are assembled. It is expected that application developers may use similar approaches to identify the minimum population size or that allow them to better quantify the expected amount of noise given a particular ε threshold.

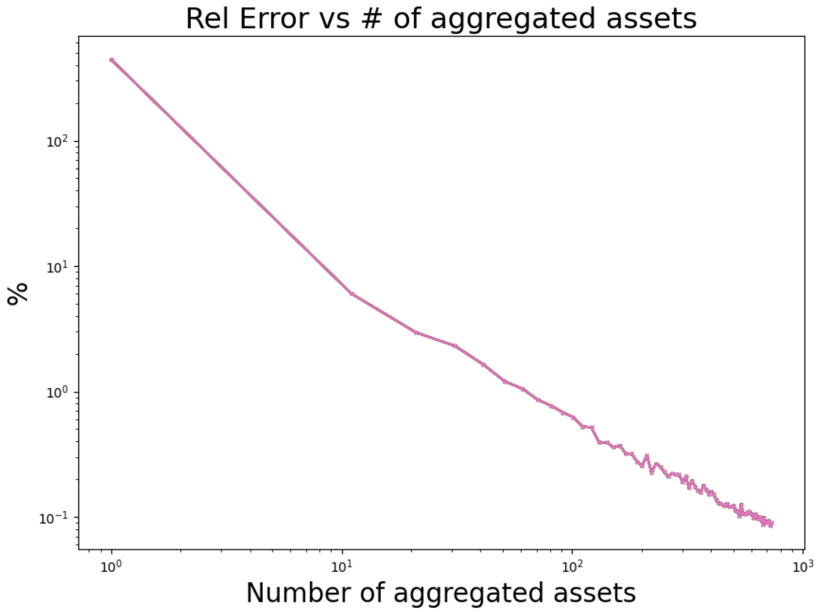

Figure 14. Relative error as a function of the number of aggregated assets. Such graphs could be used to determine the ideal aggregation size as a function of the application's error tolerance level.

3.5 Threshold Cryptography For Privacy Applications

Threshold Cryptography (TC) is a series of cryptographic schemes that issue partial keys to a set of independent actors, these partial keys can later be re-assembled to perform cryptographic functions (if the threshold conditions are satisfied). In general, a certain number of partial keys, often referred to as *shares*, must be gathered before a group key can be assembled. Assuming that n is used to indicate the number of key shares available; t is used to represent the number of shares required to reach the cryptographic threshold; and u is used to indicate the number of unique valid keys, the following conditions arise:

> $\boldsymbol{u} < \boldsymbol{t}$: No global key can be recovered; such condition protects against a malicious actor (or group thereof) that attempts to obtain access to the group key without group approval. However, this condition could deny future access if a sufficient number of agents permanently retire or lose their keys. Some implementations may allow solo individuals to verify the content of a message but prevent the complete decryption.
>
> **u>=t**: The global key can be assembled, hence enabling its use for encryption, decryption, or signature generation purposes (Vassilev, 2018). Generally speaking, a set of u independent actors have to send their partial keys to a trusted entity who then combines them together. If t is less than n, participants can continue to perform operations as long as no more than $n - t$ entities are disabled, thereby providing a level of fault tolerance (Vassilev, 2018).

As with many other cryptographic implementations, key generation and efficient key management pose one of the largest challenges for TC. As such, multiple algorithms have been proposed across the years to create and distribute key shares. For example, in the most basic implementation, a trusted entity (i.e., the dealer) generates and distributes the key shares to each participant. The participants can also rely on more advanced Distributed Key Generation (DKG) protocols to distribute the key shares among themselves if there is no trusted entity available, which may be particularly useful in DLT environments (Tomescu, et al., 2020). Naturally, multiple DKG protocols have been proposed in literature to improve upon speed, security, or to support different network architectures. For example, in (Gurkan, et al., 2021) the authors describe a gossip-based protocol to generate a pair of public and private keys that can be used to encrypt and sign content in TC-based environments. The use of gossip protocol makes it well suited for aggregating a relatively large number of key shares efficiently across asynchronous network architectures such as the internet. In addition to the various key management proposals, some researchers have sought to increase TC's basic capabilities to facilitate their integration into real-world systems. For example, in (Iftene, 2005) the authors describe a weighted-TC implementation that assigns weights to each key share. In this case, the sum of the individual weighted shares is compared against the pre-defined threshold, granting access only if $\sum k_u \geq t$. Such approaches may be leveraged to enable supervisory entities (e.g., regulators, and grid operators) to assert executive control over a group of individuals when needed.

The application of TC-based methods to increase a system's privacy characteristics is not a new subject. However, most approaches focus on applying a cryptographic function to secure a user's private data (and hence equating privacy with confidentiality). Among the works published, notable mentions include:

(Al-Muhtadi, Hill, & Al-Rwais, 2011), in which the authors proposed a multi-layer access control mechanism applicable to the healthcare industry. The proposed mechanism issues key shares to sensors, medical personnel, and other system observers who can independently evaluate if access to a medical record is justified. The multi-layer approach is created by assembling independent key holders into common groups (e.g., doctors are part of the medical personnel) and then mapping each group to a nested protection zone that essentially forces data consumers to justify their data needs at each access level.

(Cramer, Gennaro, & Schoenmakers, 1997), in which the authors describe a privacy-driven voting system that allows third parties to tally and verify the final result (via homomorphic encryption), while at the same time protecting the voting ballot's internal details. The voting details (e.g., the name and address of the voter) can only be retrieved if a sufficient number of electoral supervisors combine their data shares to reveal the encrypted payload. Such an approach could in theory be expanded to protect market bids in TES environment, allowing agents to verify the clearing price of an auction while protecting the identity of the bidder.

Based largely on the previously described constructs, a proposal for securing bidding information within DLT environments was devised. The proposed scheme relies on an out-of-band communication scheme to encrypt the transaction data (via TC) before storing it in the ledger. Due to the nature of TC, revealing the contents of the ledger (and thus the transaction itself), requires interested parties to first obtain access to the individual key shares, providing a layer of privacy that is not present in current-generation DLT implementations. The proposed scheme relies on two processes that can be daisy chained into existent applications. The secure storage process can be summarized as follows (see Figure 15 for a graphical representation):

1. Upon receiving a transaction request from the DLT network, peers validate its content (via the smart contract logic). After a successful validation, each peer generates a key share using the DKG process.

2. Due to the nature of certain DKG implementations, multiple parties may end-up with different private keys (PK_i), therefore a “trusted” merger must be elected and chosen to encrypt the transaction using its PK_i

3. The trusted merger then communicates the encrypted transaction back to the peers for them to independently reach consensus and insert the encrypted transaction into the ledger.

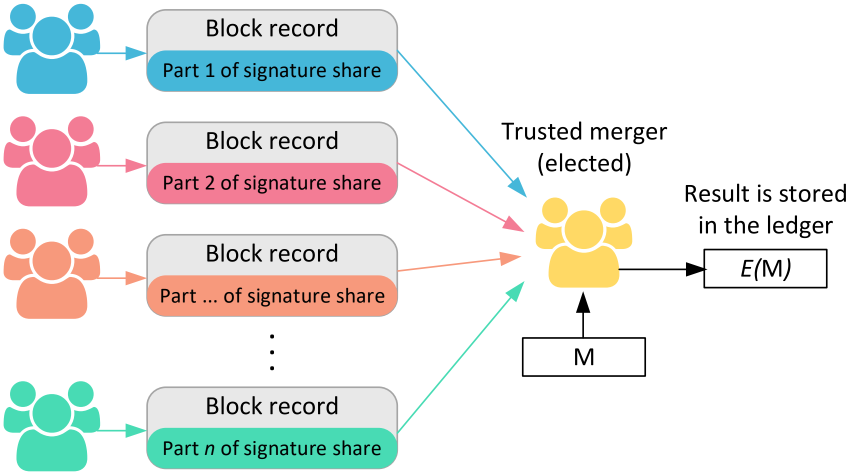
Block record
Part 1 of signature share
Block record
Part 2 of signature share
Block record
Part ... of signature share
Block record
Part *n* of signature share
Trusted merger (elected)
Result is stored in the ledger
E(M)
M

Figure 15. Overview of the proposed scheme under encryption mode

To reveal the contents of the ledger, parties must appeal to individual members and convince them to share their partial key with a trusted member of the organization. Only if enough members are convinced, then the trusted peer can decrypt and share the result with the original claim holder (see Figure 16). Due to the nature of DLT systems, the appeal process can be encoded into the smart contract to ensure that a fair approval process is followed (hence preventing bad actors from trying to block a legitimate access request).

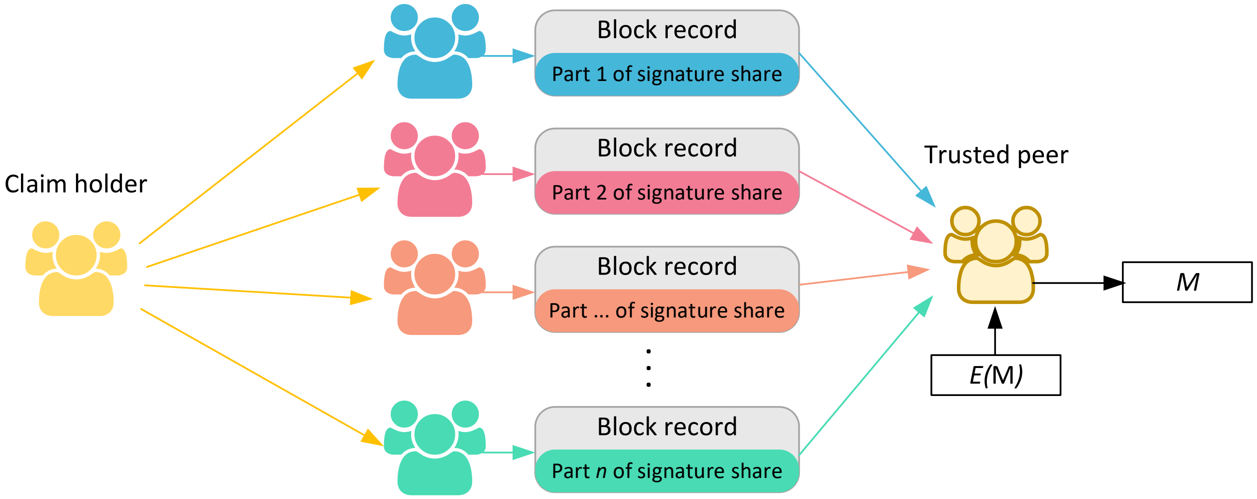
Claim holder
Block record
Part 1 of signature share
Block record
Part 2 of signature share
Block record
Part ... of signature share
Block record
Part *n* of signature share
Trusted peer
M
E(M)

Figure 16. Overview of the proposed scheme operating under decryption mode

Based on the logic described above. A proof-of-concept implementation was developed based on the Python-based library available in (Miller & Bellemare, HoneyBadgerBFT-Python, 2016). The library itself contains a fault-tolerant network deployment, named *Honey Badger,* that is intended to serve as the backbone architecture for future blockchains (Miller, Xia, Croman, Shi, & Song, 2016). *Honey Badger* is a Byzantine Fault Tolerant architecture that is designed to operate over a wide-area network under an asynchronous communication model (e.g., where network time-outs are not considered a reliable indicator of disconnection). The library implements a TC scheme that was first proposed by (Baek & Zheng, 2003), it attempts to maintain the same level of security provided by traditional TC methods but focuses on reducing the overall key length. The use of shorter key lengths can help improve the performance of distributed systems by requiring smaller packet sizes that are easier to transmit. However, *Honey Badger* is not a full-fledged blockchain implementation and thus cannot support the use of Smart contracts (SC) and as such, only the cryptographic functions could be tested.

In addition to the aforementioned limitations, the use of python as a smart contract language is not yet supported by the Hyperledger Fabric project (which is the implementation used by the team in the past to evaluate TES market approaches). Therefore, at the current stage, the proposed approach remains an architectural concept that awaits the maturation of libraries, and other supporting infrastructure before it can be assembled and tested in a lab-scale deployment (see Figure 17).

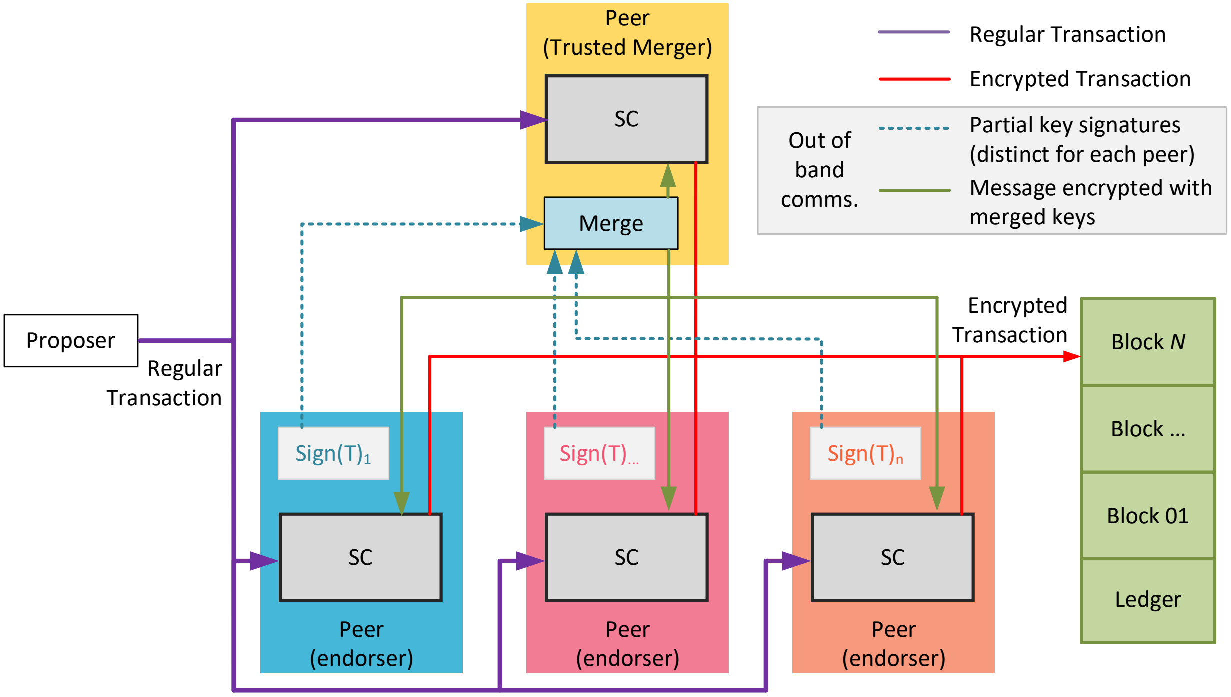

Figure 17. Data exchanges of the proposed architecture under encryption mode.

4 Recommended Future Work

The electrical grid is often described as one of the largest, and most complex machines that humans have ever built. Across the decades the system has undergone several technological transformations to address emerging needs, often resulting in a series of incremental updates that are built upon each other. However, addressing many of the outstanding electrification and resilience goals will require a much more transformative approach to address the structural limitations of the current grid. This will undoubtedly require the development of new (or improved) operational schemes and associated control architectures that will better prepare the grid for the future. Despite the challenges, the ongoing transformation offers system architects the opportunity to build systems that are more inclusive of the needs and expectations of their participants. As outlined by this work, one important aspect to consider is data privacy, which if done right, could enable the emergence of new operational use cases that allow individuals to assemble and participate in communities to provide, exchange, or benefit from grid-related services.

The successful adoption of privacy aspects in energy applications will require a coordinated effort between solution developers, academia, and public advocates to build solutions that not only satisfy the application-specific needs but also offer privacy guarantees that serve as a welcoming feature for future users. From an academic perspective, privacy efforts within the energy space are likely to fall into one or more of the following tracks: a) Developing abstract technologies or methods that can be used to enhance privacy, b) Identifying systematic gaps and evaluating potential solutions applicable to the energy space and c) Creating application-specific solutions that seek to integrate privacy as a design pillar. These tracks, although mostly independent from each other, could build upon each other to increase the overall maturity levels of the field. Based on the literature review discussed in Section 2 and the observed limitations, research tracks that may be worth pursuing in the near future may include:

- Continuing the research and development of privacy mechanisms. Such mechanisms could be used to enforce privacy guarantees that enable data owners to share data openly and securely with third parties while minimizing the risk of inadvertently exposing personal or business-sensitive data. Although individual solutions are likely to be application-specific, the basic constructs should be generalizable to a wide set of problems (e.g., to protect highly repetitive time series data, or to produce deanonymized topological models)

- Developing methodologies that can be used to identify and evaluate existing data production and consumption processes. The obtained results could then be used to gradually move towards a *"need-driven"* data usage model. For example, data collection processes should seek to minimize the production of excessive or duplicate data records to minimize future risks (e.g., data leaks). This may include developing frameworks or automated tools that enable the industry to independently assess their privacy posture, such tools may build and adapt from related frameworks applicable to the cybersecurity space.

- Creating industry best practices and related tools that can serve as guidelines for developing and implementing secure, privacy-aware solutions. Such tools could build upon cybersecurity principles yet must remain conscious that privacy needs to be a user-driven process. Therefore, policies for the collection, use, and eventual destruction of data records must be tailored to the individuals, which may themselves establish different sensitivity levels that must be accounted for during the design phase.

- Developing a set of standardized data sets and associated benchmarks that allow researchers to test and compare the effectiveness of the proposed privacy solutions among their peers. In the past, such repositories have enabled the rapid advancement of knowledge within a particular domain. Examples include the image repositories used to assess image compression technology and the use of labeled image libraries to evaluate machine learning algorithms' performance.

- Increasing public awareness towards data privacy. This may include fostering risk awareness, but also identifying the benefits of implementing secure data sharing to improve user experience, enabling new operational use cases to be implemented. This may be particularly useful in transitioning away from the “restrict by default” policy that most energy sector organizations follow into a “managed risk sharing” process that allows industry and academia to increase their collaboration.

Clearly, the aforementioned research directions could result in the successful adoption of privacy as a design pillar of future applications. At the same time, it’s imperative that organizations and solution developers are aware of potential pitfalls that could lead to a false sense of security, these may include:

- Not considering the effects of releasing overlapping data sets. The existence of overlapping datasets can enable attackers to apply set operators (i.e., union, intersection, complement) thereby helping them to increase their knowledge about individual dataset members. Theorems such as the one described in (Dwork & Roth, 2014) could be used to estimate the number of overlapping, noise-protected datasets that may lead to a comprise. In case online query engines are implemented, special care must be put into ensuring the supported queries do not allow an external agent to assemble overlapping datasets that may break the privacy protections (via a differentiating attack).

- Not realizing the lifespan of their products. Solutions developed today may satisfy current privacy requirements but may fail to do so in a long-term scenario if the methods become vulnerable or too much data is eventually released. This is particularly concerning in the energy space, as grid applications are expected to remain operational for long periods of time (e.g., decades). Note that disclosing older datasets may be acceptable, but their release must be a planned event rather than an unexpected result.

- Failing to evaluate a method under the specifics of the use case or implementation. Methods that are considered secure do not necessarily yield secure implementations, this may be due to deficient coding practices, but also to inherent features of the use case that introduce vulnerabilities. An example of such a scenario was described in Section 2, where side-channel attacks were discussed.

5 Conclusions

This report presented a high-level overview of digital privacy in the context of the power grid. Digital privacy by itself is a complex topic that touches on a variety of subjective and technical aspects that extend over the legal, social, and engineering domains, but whose main goal is to enable individuals to assert control over their digital footprint. At the same time, the grid continues (and will likely continue) to embrace a communications-driven approach that generates vast amounts of data which makes it vulnerable to potential abuse by ill-intended entities. Based on the current landscape, this work presented an overview of grid-related privacy threats documented in the literature, as well as summarizing techniques that may help to improve the privacy posture of grid applications.

This report highlights the need for privacy considerations during the design, implementation, and eventual operation of grid-related applications. The adoption of privacy as a design pillar could serve as an enabler for data-driven solutions where actors can openly collaborate with peers while maintaining their individual behaviors private from other members. Privacy-aware environments could provide the necessary technical constructs and guarantees needed to support the collaboration of DERs, PVs, service aggregators, and other entities in a fair and inclusive environment. Furthermore, such mechanisms could allow grid operators to maintain a supervisory role to ensure grid safety and operational constraints are maintained despite the shift from a centralized to a decentralized grid environment.

This work made a series of recommendations to encourage system designers to rethink the current practices, which typically rely on the mass collection of data into a targeted, need-driven data collection solution approach that minimizes privacy risks. In addition, it presented a technical evaluation of Differential Privacy (DP) within the context of Transactive Energy Systems. The results indicate that DP-based mechanisms, like many other constructs, are technically sound but may underperform due to poor implementations, or when mismatches between the method’s strengths and the application’s environment exist. Therefore, it is imperative that technologies are screened and evaluated to ensure privacy goals are met, such evaluations should be cognizant of the relatively long-term lifecycles, and the vast amounts of data that energy applications require.

In summary, this work highlights the need for 1) Supporting the development of technical constructs that can be used to protect and enforce privacy expectations; 2) Raising awareness and creating methodologies that allow system developers to integrate privacy into their designs; and 3) Developing solutions that can achieve functional goals while remaining privacy aware.

6 References

Almadhoun, N., Ayday, E., & Ulusoy, Ö. (2020). Inference attacks against differentially private query results from genomic datasets including dependent tuples. *Bioinformatics*, i136–i145.

Al-Muhtadi, J., Hill, R., & Al-Rwais, S. (2011). Access control using threshold cryptography for ubiquitous computing environments. *Journal of King Saud University-Computer and Information Sciences*.

Apthorpe, N. D. (2017). Spying on the Smart Home: Privacy Attacks and Defenses on Encrypted IoT Traffic. *arXiv preprint arXiv:1708.05044*.

Atluri, G., Karpatne, A., & Kumar, V. (2018). Spatio-Temporal Data Mining: A Survey of Problems and Methods. *ACM Computing Surveys*.

Badr, M. M. (2022). Privacy-Preserving Federated-Learning-Based Net-Energy Forecasting. *SoutheastCon 2022* (pp. 133-139). Mobile: IEEE.

Badr, M. M. (2023). Privacy-Preserving and Communication-Efficient Energy Prediction Scheme Based on Federated Learning for Smart Grids. *IEEE Internet of Things Journal*, 7719-7736.

Baek, J., & Zheng, Y. (2003). Simple and efficient threshold cryptosystem from the Gap Diffie-Hellman group. *IEEE Global Telecommunications Conference.*

Braun, T. B. (2018). Security and Privacy Challenges in Smart Cities. *Sustainable Cities and Society*, 499-507.

Brighente, A. M. (2023). Electric Vehicles Security and Privacy: Challenges, Solutions, and Future Needs. *arXiv preprint arXiv:2301.04587*.

Chen, D. a. (2017). Weatherman: Exposing Weather-Based Privacy Threats in Big Energy Data. *2017 IEEE International Conference on Big Data (Big Data)* (pp. 1079-1086). Boston: IEEE.

Chen, D. S. (2016). Sunspot: Exposing the Location of Anonymous Solar-Powered Homes. *The 3rd ACM International Conference on Systems for Energy-Efficient Built Environments* (pp. 85-94). Palo Alto: Association for Computing Machinery.

Cramer, R., Gennaro, R., & Schoenmakers, B. (1997). A secure and optimally efficient multi-authority election scheme. *European transactions on Telecommunications*.

Darby, Sarah J. (2012). Metering: EU policy and implications for fuel poor households. *Energy Policy*.

Dasom, L., & Hess, D. (2021). Data privacy and residential smart meters: Comparative analysis and harmonization potential. *Utilities Policy*.

di Vimercati, S. D. (2023). K-Anonymity: From Theory to Applications. *Transactions on Data Privacy*, 25-49.

Dwork, C., & Roth, A. (2014). Ch 8. Lower Bounds and Separation Results. In C. Dwork, & A. Roth, *The Algorithmic Foundations of Differential Privacy* (pp. 158-170).

El Emam, K., Jonker, E., Arbuckle, L., & Malin, B. (2011). A Systematic Review of Re-Identification Attacks on Health Data. *PLOS ONE*.

Fan, J. Q. (2017). Privacy Disclosure Through Smart Meters: Reactive Power Based Attack and Defense. *2017 47th Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN)* (pp. 13-24). Denver: IEEE.

Faustine, A., Mvungi, N. H., Kaijage, S., & Kisangiri, M. (2017). A Survey on Non-Intrusive Load Monitoring Methodies and Techniques for Energy Disaggregation Problem.

Fraiji, Y. L. (2018). Cyber Security Issues of Internet of Electric Vehicles. *2018 IEEE Wireless Communications and Networking Conference (WCNC)* (pp. 1-6). Barcelona: IEEE.

Gourisetti, S. N., Sebastian-Cardenas, D. J., Bhattarai, B., Wang, P., Widergren, S., Borkum, M., & Randall, A. (2021). Blockchain Smart Contract Reference Framework and Program Logic Architecture for Transactive Energy Systems. *Applied Energy*, 0306-2619.

GridWise Architecture Council. (2018). *GridWise Transactive Energy Framework.* PNNL.

Gurkan, K., Jovanovic, P., Maller, M., Meiklejohn, S., Stern, G., & Tomescu, A. (2021). Gurkan, Kobi and Jovanovic, Philipp and Maller, Mary and Meiklejohn, Sarah and Stern, Gilad and Tomescu, Alin. *Annual International Conference on the Theory and Applications of Cryptographic Techniques.*

Han, W. a. (2016). Privacy Preservation for V2G Networks in Smart Grid: A Survey. *Computer Communications*, 17-28.

Hawes, M. (2021). *Understanding the 2020 Census Disclosure Avoidance System: Simulated Reconstruction-Abetted Re-identification Attack on the 2010 Census .* U.S. Census Bureau.

Hoh, B. M. (2006). Enhancing Security and Privacy in Traffic-Monitoring Systems. *IEEE Pervasive Computing*, 38-46.

Huang, C. M. (2017). Preserving Source Location Privacy for Energy Harvesting WSNs. *Sensors*, 724.

Iftene, S. a. (2005). Weighted Threshold Secret Sharing Based on the Chinese Remainder Theorem. *Scientific Annals of Cuza University*, 161-172.

JackBarber98. (2019). *pruned-nilm*. Retrieved from https://github.com/JackBarber98/pruned-nilm

Kiarie, L. K. (2019). Application of Spritz Encryption in Smart Meters to Protect Consumer Data. *Journal of Computer Networks and Communications*.

Kimani, K. V. (2019). Cyber Security Challenges for IoT-Based Smart Grid Networks. *International Journal of Critical Infrastructure Protection*, 36-49.

Kumar, G. R.-K. (2020). A Privacy-Preserving Secure Framework for Electric Vehicles in IoT Using Matching Market and Signcryption. *IEEE Transactions on Vehicular Technology*, 7707-7722.

Laurens, R., Christianto, E., Caulkins, B., & Zou, C. C. (n.d.). Side-Channel VoIP Profiling Attack against Customer Service Automated Phone System. *GLOBECOM 2022 - 2022 IEEE Global Communications Conference*, (pp. 6091-6096).

Lee, J., & Clifton, C. (2011). How Much Is Enough? Choosing ε for Differential Privacy. *Information Security Conference*.

Lei, Y.-T. C.-Q.-S.-Q. (2022). A Renewable Energy Microgrids Trading Management Platform Based on Permissioned Blockchain. *Energy Economics*, 106375.

Lets Make It LLC. (2020). *Differential Privacy: An NPM library*. Retrieved from https://github.com/gittyeric/differential-privacy

Li, W. H.-C. (2020). Privacy-Preserving Traffic Management: A Blockchain and Zero-Knowledge Proof Inspired Approach. *IEEE access*, 181733-181743.

Lindholm, M., Richman, R. T., & M.V., W. (2022). DISCRIMINATION-FREE INSURANCE PRICING. *ASTIN Bulletin: The Journal of the IAA*.

Liu, C., Chakraborty, S., & Mittal, P. (2016). Dependence Makes You Vulnerable: Differential Privacy Under Dependent Tuples. *Network and Distributed System Security (NDSS) Symposium.* San Diego, CA.

Liyue, F., & Li, X. (2014). An Adaptive Approach to Real-Time Aggregate Monitoring With Differential Privacy. *IEEE Transactions on Knowledge and Data Engineering*.

Long, G. (2020). *Formal Privacy Methods for the 2020 Census.* McLean, Virginia. Retrieved from https://www2.census.gov/programs-surveys/decennial/2020/program-management/planning-docs/privacy-methods-2020-census.pdf

Mahmood, K. S. (2018). An elliptic curve cryptography based lightweight authentication scheme for smart grid communication. *Future Generation Computer Systems*, 557-565.

Miller, A., & Bellemare, S. (2016). *HoneyBadgerBFT-Python*. Retrieved from https://github.com/initc3/HoneyBadgerBFT-Python/tree/dev

Miller, A., Xia, Y., Croman, K., Shi, E., & Song, D. (2016). The honey badger of BFT protocols. *ACM SIGSAC conference on computer and communications security.*

Mironov, I. (2017). Rényi Differential Privacy,. *IEEE 30th Computer Security Foundations Symposium (CSF).*

Mivule, K. (2013). Utilizing Noise Addition for Data Privacy, an Overview. *International Conference on Information and Knowledge Engineering .*

Narayanan, A., & Shmatikov, V. (2008). Robust De-anonymization of Large Sparse Datasets. *2008 IEEE Symposium on Security and Privacy*.

Nayak, T. (2020). *A Review of Rigorous Randomized Response Methods for Protecting Respondent's Privacy and Data Confidentiality.* Center for Statistical Research and Methodology, U.S. Census Bureau.

Northwest Energy Efficiency Alliance. (2020). *Home Energy Metering Study: Public Data User Guide.* Portland, Oregon.

Oak Ridge National Laboratory. (2023, January 26). *Vandy Tombs – Using Math to Improve Data Privacy in Machine Learning*. Retrieved October 16, 2023, from https://www.ornl.gov/organization-news/vandy-tombs-using-math-improve-data-privacy-machine-learning

Pal, R. P. (2018). Privacy Engineering for the Smart Micro-Grid. *IEEE Transactions on Knowledge and Data Engineering*, 965-980.

Pepermans, G. (2014). Valuing smart meters. *Energy Economics*, 280-294.

Poh, G. S. (2019). PrivHome: Privacy-Preserving Authenticated Communication in Smart Home Environment. *IEEE Transactions on Dependable and Secure Computing*, 1095-1107.

Rabieh, K. M. (2016). Privacy-Preserving Route Reporting Schemes for Traffic Management Systems. *IEEE Transactions on Vehicular Technology*, 2703-2713.

Razavi, R. A. (2019). Occupancy Detection of Residential Buildings Using Smart Meter Data: A Large-Scale Study. *Energy and Buildings*, 195-208.

Shokri, R. M. (2017). Membership Inference Attacks Against Machine Learning Models. *2017 IEEE symposium on security and privacy (SP)* (pp. 3-18). San Jose: IEEE.

Shuaib, K. Z.-H. (2015). Resiliency of Smart Power Meters to Common Security Attacks. *Procedia Computer Science*, 145-152.

Son, Y.-B. J.-H.-Y.-Y.-K. (2020). Privacy-Preserving Peer-to-Peer Energy Trading in Blockchain-Enabled Smart Grids Using Functional Encryption. *Energies*, 1321.

Strepparava, D. F. (2022). Privacy and Auditability in the Local Energy Market of an Energy Community with Homomorphic Encryption. *Energies*, 5386.

Sun, X. F. (2021). A Survey on Zero-Knowledge Proof in Blockchain. *IEEE network*, 198-205.

Teng, F., Chhachhi, S., Ge, P., Graham, J., & Gunduz, D. (2022). *Balancing Privacy and Access to Smart Meter Data.* London: Energy futures Lab, An institute of Imperial College London.

Tomescu, A., Chen, R., Zheng, Y., Abraham, I., Pinkas, B., Gueta, G. G., & Devadas, S. (2020). Towards scalable threshold cryptosystems. *IEEE Symposium on Security and Privacy.*

Tramer, F., Boneh, D., & Paterson, K. (2020). Remote Side-channels Attacks on Anonymous Transactions. *29th USENIX security symposium*.

Tran, H.-Y. J. (2022). Smart meter data obfuscation with a hybrid privacy-preserving data publishing scheme without a trusted third party. *IEEE Internet of Things Journal*, 16080-16095.

Vassilev, A. N. (2018). Psst, Can You Keep a Secret? *Computer*, 94-97.

Wang, K.-C. Y. (2021). Variational Model Inversion Attacks. *Advances in Neural Information Processing Systems*, 9706-9719.

Warner, S. L. (1965). Randomized response: a survey technique for eliminating evasive answer bias. *J Am Stat Assoc.*

Wright, C. V., Ballard, L., Coull, S. E., Monrose, F., & Masson, G. M. (2008). Spot Me if You Can: Uncovering Spoken Phrases in Encrypted VoIP Conversations. *IEEE Symposium on Security and Privacy* , pp. 35-49.

Xu, H. W. (2017). CTS-DP: Publishing correlated time-series data via Differential Privacy. *Knowledge-Based Systems*.

Xu, S. X. (2021). EVchain: an anonymous blockchain-based system for charging-connected electric vehicles. *Tsinghua Science and Technology*, 845-856.

Yin, L. J. (2021). A Privacy-Preserving Federated Learning for Multiparty Data Sharing in Social IoTs. *IEEE Transactions on Network Science and Engineering*, 2706-2718.

Zeng, X. Q. (2017). A Lightweight Privacy-Preserving Scheme for Metering Data Collection in Smart Grid. *2017 IEEE 18th International Symposium on A World of Wireless, Mobile and Multimedia Networks (WoWMoM)* (pp. 1-6). Macau: IEEE.

Zhang, C. Y. (2021). A Survey on Federated Learning. *Knowledge-Based Systems*, 106775.

Zhang, C., Zhong, M., Wang, Z., Goddard, N., & Sutton, C. (n.d.). Sequence-to-point learning with neural networks for nonintrusive load monitoring.

Zhang, Y. J. (2018). Privacy-Preserving Communication and Power Injection Over Vehicle Networks and 5G Smart Grid Slice. *Journal of Network and Computer Applications*, 50-60.

Zhu, Y., Lu, Y., & Vikram, A. (2012). On Privacy of Encrypted Speech Communications. *IEEE Transactions on Dependable and Secure Computing*.